A Generic Truthful Mechanism for Combinatorial Auctions

Hanrui Zhang

Duke University, Durham NC 27705, USA
hrzhang@cs.duke.edu

Abstract. We study combinatorial auctions with n agents and m items, where the goal is to allocate the items to the agents such that the social welfare is maximized. We present a universally truthful mechanism with polynomially many queries for combinatorial auctions. Our mechanism and analysis work adaptively for all classes of valuation functions, guaranteeing $\widetilde{O}(\min(d, \sqrt{m}))$-approximation of the optimal social welfare, where d is the degree of complementarity of the valuation functions. To our knowledge, this is the first mechanism that achieves an approximation guarantee better than $\Omega(\sqrt{m})$, when the valuations exhibit any kind of complementarity.

Keywords: Truthful combinatorial auctions, Approximate subadditivity, Pointwise approximation

1 Introduction

The field of algorithmic mechanism design studies protocols for computing an outcome to optimize a certain social objective (e.g., the social welfare), when inputs are reported by strategic agents. The main challenge in algorithmic mechanism design is twofold: algorithmically, the mechanism has to deal with the *computational hardness* of the problem; strategically, the mechanism has to take into account the *incentives* of the agents, which often do not align with the interests of the designer. One popular scheme in the field is to design *truthful* mechanisms, where the dominant strategy of all bidders are to report their true preferences. Restricted to truthful mechanisms, one no longer needs to worry about complex strategic behavior, and can therefore focus on the algorithmic properties of the mechanism.

In this paper, we consider a central problem in algorithmic mechanism design — designing truthful mechanisms for *combinatorial auctions*. In a combinatorial auction, there are n *agents* and m *items*. Each agent i has a *valuation function* v_i, that maps each subset S of the items to her value of the subset $v_i(S)$. The goal is to find an *allocation* of all items, $(A_1, \dots, A_n)$, such that the total value (i.e., the *social welfare*) of the agents, $\sum_{i \in [n]} v_i(A_i)$, is maximized. It is standard in combinatorial auctions to assume that all valuations are *monotone*[1]

[1] A valuation v is monotone, if for any $S \subseteq T \subseteq [m]$, $v(S) \leq v(T)$.

and *normalized*[2]. Previous research also studies restricted classes of valuations, e.g., *submodular*[3], *fractionally subadditive (XOS)*[4], and *subadditive*[5] valuations. It is known that all submodular valuations are fractionally subadditive, and all fractionally subadditive valuations are subadditive.

Since the size of a valuation function can be exponentially large in m, it is often impossible to use the entire functions as the input. Instead, two standard kinds of queries are allowed: (1) *value* queries, which, given an agent i and a set S, return the value of S to agent i, $v_i(S)$; (2) *demand* queries, which, given an agent i and prices $\{p_j\}_{j\in[m]}$, return a utility-maximizing set (i.e., a *demand set*) of i under the given prices. That is, the query returns a set S that maximizes $v_i(S) - \sum_{j\in S} p_j$.

Combinatorial auctions become relatively easy if we remove either one of the two aspects of the difficulty. Ignoring incentive issues, efficient approximation algorithms exist for the welfare maximization problem. Vondrak gives a $\frac{e}{e-1}$-approximation for submodular valuations, using only value queries [26], which is shown tight by Mirrokni et al. [23]. When demand queries are allowed, Feige and Vondrak give an upper bound of $\frac{e}{e-1} - 10^{-6}$ for submodular valuations [17], where a lower bound of $\frac{2e}{2e-1}$ is known [10]. Feige gives a $\frac{e}{e-1}$-approximation for XOS valuations and a 2-approximation for subadditive valuations using both queries [14]. None of these algorithms are truthful. On the other hand, the VCG mechanism is truthful and guarantees the optimal welfare. Computing the VCG outcome and payments, however, is usually algorithmically hard. In particular, approximation usually does not help in implementing the mechanism because of incentive issues.

Taking into account both computational and strategic issues, there are significant gaps between known upper and lower bounds. Under the most restrictive assumptions, for submodular valuations, Dobzinski et al. [8] give a deterministic $O(\sqrt{m})$-approximation that requires only value queries, which is tight both information theoretically [6] and complexity theoretically [10]. Allowing randomization and demand queries, a series of work improves the upper bound from $O(\log^2 m)$ for XOS valuations [9], to $O(\log m \log\log m)$ for subadditive valuations [5], to $O(\log m)$ for XOS valuations [21], to $O(\sqrt{\log m})$ for XOS valuations [7], and finally to $O((\log\log m)^3)$ for XOS valuations [2]. For general valuations, $O(\sqrt{m})$-approximation randomized mechanisms using both kinds of queries are known [9, 5], accompanied by a matching $\Omega(m^{1/2-\varepsilon})$ communication complexity lower bound by Nisan [24].

All of the above mechanisms are *universally truthful*. That is, fixing the randomness of the mechanism, no agent has incentive to misreport her valuation. We focus our attention on universally truthful mechanisms, as opposed to *truthful in expectation* ones, since if the mechanism proceeds in stages, as agents observe

[2] A valuation v is normalized, if $v(\emptyset) = 0$.

[3] A valuation v is submodular, if for any $S, T \subseteq [m]$, $v(S)+v(T) \geq v(S\cup T)+v(S\cap T)$.

[4] A valuation v is fractionally subadditive, if for any S, $\{T_i\}$, and $\{\alpha_i\}$, $v(S) \leq \sum \alpha_i v(T_i)$, whenever the following holds: for each $j \in S$, $\sum_{i:j\in T_i} \alpha_i \geq 1$.

[5] A valuation v is subadditive, if for any $S, T \subseteq [m]$, $v(S) + v(T) \geq v(S \cup T)$.

partial realization of the randomness, truthfulness in expectation may not be able to prevent them from lying. Even if agents do not observe the realization of the randomness, their attitude toward risk may still lead them to misreport.

Despite all the upper bounds for various restricted classes of valuations, little is known for classes beyond subadditivity. Subadditive valuations are considered reasonably general, but they can only model items as *substitutes* to each other — that is, possessing some items can never make other items more desirable. While focusing on subadditive valuations usually allows better approximation ratios, real world valuations often do involve *complementarity*. For example, a TV set seems more valuable when one already has a sofa, because otherwise she might have to watch on her feet. On the other hand, the amount of complementarity is usually *limited*, in the sense that a sofa and a TV set complement each other, but neither of them would affect the value of a car, a dishwasher, or anything out of the living room. In other words, possible sets of items that complement each other are likely not too large. Such valuations with limited complementarity, while being obviously more general than the subadditive class, still seem intuitively easier to handle than arbitrary monotone valuations. So, a natural question arises:

> ***Beyond subadditivity, can we do better than $\Omega(\sqrt{m})$, when agents have valuations exhibiting limited complementarity?***

1.1 Our Results

We give a positive answer to the question above. Our main contribution is twofold:

1. Going beyond subadditive valuations, we establish welfare guarantees that degrade smoothly as the degree of complementarity of the valuations grows. We prove fine-grained upper bounds roughly proportional to the degree of complementarity, which, when the degree is small, improve substantially over the $O(\sqrt{m})$ bound for general valuations. To our knowledge, no such results were known before.
2. We provide unified design and analysis that work adaptively for all classes of valuations, guaranteeing approximation ratios that nearly match the state-of-the-art for the respective class.

In order to derive parametrized welfare guarantees for valuations beyond the complement-free class, we need to be able to measure how much complementarity the valuations exhibit (i.e., we need a *measure of complementarity*). While several measures have been proposed and referred to in various applications (e.g., the supermodular degree hierarchy [16] and the Maximum-over-Positive-Hypergraphs hierarchy [15]), it has been observed that different tasks often require different measures to capture the transition of hardness from restricted to general valuations (see, e.g., [12]). For our problem, the superadditive width hierarchy proposed by Chen et al. [3] seems the best fit. The measure builds on the concept of superadditive sets:

Definition 1 (Superadditive Sets [3]). *Let $v(S|T) = v(S \cup T) - v(T)$ be the marginal of S given T. Given a normalized monotone valuation function v over a ground set $[m]$, a set $T \subseteq [m]$ is* superadditive *w.r.t. v if*

$$\exists\, S \subseteq [m] \setminus T \text{ such that: } v(S|T) > \max_{T' \subsetneq T} v(S|T').$$

In words, a set T is superadditive, if it enables some set S with a larger marginal than any of its proper subsets does. Based on the concept of superadditive sets, Chen et al. define a measure of complementarity:

Definition 2 (Superadditive Widths [3]). *The* superadditive width *of a valuation function v is defined to be*

$$\mathrm{SAW}(v) = \max\{|T| \mid T \text{ is a superadditive set w.r.t. } v\}.$$

The definition essentially says, that the degree of complementarity of a valuation is proportional to the size of the largest superadditive set with respect to the valuation.

It is known that for any monotone valuation function v over $2^{[m]}$, $0 \leq \mathrm{SAW}(v) \leq m-1$, and $\mathrm{SAW}(v) = 0$ iff v is subadditive [3]. In other words, valuations can be categorized, according to their superadditive width, into m nested layers, where the lowest layer (layer 0) contains exactly the class of subadditive valuations, and the highest layer (layer $m-1$) contains all monotone valuation functions. We denote the d-th layer, containing valuations with superadditive width at most d, by SAW-d.

The following theorem summarizes our results:

Theorem 1 (Informal). *There is an efficient universally truthful mechanism which guarantees $\widetilde{O}(\min(d, \sqrt{m}))$-approximation*[6] *of the optimal welfare, where m is the number of items, and $d = \max_{i \in [n]} \mathrm{SAW}(v_i)$ is the maximum superadditive width of agents' valuations.*

	Submodular/XOS	Subadditive	SAW-d	General
Mechanism 1 of [5]	$O(\sqrt{m})$	$O(\sqrt{m})$	$O(\sqrt{m})$	$O(\sqrt{m})$
Mechanism 2 of [5]	$O(\log m \log\log m)$	$O(\log m \log\log m)$	?	?
[2]	$O((\log\log m)^3)$	?	?	?
This paper	$O(\log m)$	$O(\log^2 m)$	$O(d \log^2 m)$	$O(\sqrt{m} \log m)$

Table 1. Comparison of approximation ratios of several mechanisms.

The mechanism and analysis we present enjoy generic applicability — they require no parameters and automatically work for all kinds of valuations. Beside our result for limited-complementarity valuations, for complement-free valuations, we recover the polylog approximation ratios, and for general valuations,

[6] $\widetilde{O}$ hides a polylog m factor.

we match the $\Omega(\sqrt{m})$ lower bound up to a $O(\sqrt{\log m})$ factor. This adaptivity is particularly desirable when it is unrealistic to know beforehand to which class the valuations belong[7]. We also note that our mechanism is considerably simplified compared to previously proposed mechanisms — we intend to keep the mechanism as simple as possible to demonstrate the power of the underlying ideas, potentially compromising a minor factor in the approximation ratio. On the other hand, our analysis does shed light on the potential space for improvement within the framework we present, possibly by incorporating ideas from [5, 2]. For further related work, see Appendix A.

1.2 Organization and Technical Overview

We present our mechanism in Section 2, and then proceed to establishing approximation guarantees for different classes of valuations in later sections. The overall idea is to build a framework using the strongest assumptions under which the argument remains illustrative, and then generalize gradually by adapting the framework.

We begin our investigation with *constraint homogeneous (CH)* valuations (defined in Definition 3), which is arguably the simplest class of valuations exhibiting complementarity. The class was originally introduced by Devanur et al. [4] and extended by Feldman et al. [18] to study the PoA of simple auctions. Roughly speaking, the CH class contains valuations that are additive over small disjoint bundles, where each bundle's value is proportional to its size. We show in Section 3 that our mechanism guarantees $\widetilde{O}(d)$-approximation for CH valuations with maximum bundle size d. More specifically, we first show that given complete information about agents' valuations, there exist prices, such that if we post these prices on the items, order agents arbitrarily, and let them purchase their demand sets, the resulting allocation is a $O(d)$-approximation of the optimal welfare. We prove this guarantee using a standard argument that decomposes the welfare into two parts: the total payment, and the total buyer surplus. The intuition is that, if we post the right prices, then when most items are sold, the payment must be high enough. Otherwise, since the unsold items are available to every agent as an option, the total buyer surplus must be high enough. The welfare bound follows since both terms are nonnegative. We then argue that without knowing agents' valuations, we can somehow guess a price, such that if we post that price on every item, the expected welfare is still reasonably high. The technique of "guessing a price" has also been shown useful in [9, 5].

We further observe that for certain truthful mechanisms, *pointwise approximation* between classes of valuations (as defined in Definition 4) in a sense

[7] One may argue that running the state-of-the-art mechanism for each class of valuations with constant probability achieves the best approximation guarantee for all classes simultaneously. The point we try to make here is, we show how one can achieve this adaptivity with coherent design and analysis, which arguably provides more insight into the problem, and is more likely to inspire future research on the topic.

preserves welfare guarantees. The notion of pointwise approximation was explicitly defined by Devanur et al. in [4], where they show such approximation approximately preserves PoA bounds. Informally, v is approximated by v' at set S, if (1) v' is always no larger than v at any subset of S, and (2) v' is not too much smaller than v at S. In Section 4, based on this observation, we provide a way to translate these approximation relationships into welfare guarantees, by proving the following lemma:

Lemma 1 (Informal). *There is an efficient universally truthful mechanism which guarantees $\widetilde{O}(d)$-approximation of the optimal welfare, when agents have valuations approximated by disjoint bundle (DB) valuations (as defined in Definition 5) with maximum bundle size d.*

The class of DB valuations is similar to CH, except that each bundle can have an arbitrary value. We first argue the lemma for CH valuations, and then extend to DB valuations by assigning a dummy agent to every bundle in a DB valuation. The proof of the lemma builds on the observation, that if we pretend that the agents have CH valuations that approximate the actual ones at some optimal allocation, we can borrow the argument for CH valuations with local modifications. In particular, since the welfare of the optimal allocation under the dummy valuations is not too much smaller than the actual optimal welfare, we can use the dummy welfare as the benchmark without significant loss.

The extension lemma above essentially says, in order to establish welfare guarantee for a particular class of valuations, one only needs to show approximability of the class by DB valuations. Given the extension lemma, we plug in previously known approximation results for XOS, subadditive, and SAW-d valuations by DB valuations, which immediately yields approximation guarantees for the respective classes of valuations.

Finally, in Section 5, we show that for general valuations, we are able to nearly recover the optimal $O(\sqrt{m})$ approximation ratio. We take a similar but slightly different approach. We argue that if the agents' shares in the optimal allocation are roughly equally distributed, then we can ignore agents who receive too many items. The intuition is, since agents receive disjoint sets of items, the number of agents who receive many items is not too large. Also, since the optimal welfare is equally distributed, a small number of agents cannot share too large a fraction of the welfare, and can therefore be removed without hurting the welfare too much. We then use the optimal allocation projected to agents who receive few items as the benchmark. We observe, that the valuation of each agent is approximated at the set she receives, by a CH valuation with reasonably small maximum bundle size. A similar argument to the one we use to prove the extension lemma gives the desired approximation guarantee.

2 A Generic Mechanism

In this section, we present our generic mechanism for truthful combinatorial auctions, and state its approximation guarantees for different classes of valuations.

Notation. Throughout the paper we use $[n]$ and $[m]$ to denote the sets of agents and items, respectively. W.l.o.g. we assume $m = 2^p$ for some integer p. In general we use i as the index of an agent, and j the index of an item.

The mechanism, as well as the frameworks presented in [9, 5], uses two widely applied subroutines:

– A (grand-bundle) *second price auction*, where each agent bids on the grand bundle of all items. The agent with the highest bid wins, receives all items, and pays the second highest bid.
– A *fixed-price auction* with price p, where all agents are approached in some arbitrary order. Each agent, when being asked, can choose to purchase any subset of the items available at the time, paying p for each item she purchases. Any item purchased by some agent becomes unavailable immediately.

A generic mechanism.

1. With probability $\frac{1}{2}$, run a second price auction on the grand bundle, give all items to the winner, charge her the second highest bid, and terminate.
2. Partition all agents into two sets: STAT and FIXED. Each bidder is assigned independently, with probability $\frac{1}{2}$ to STAT, and with probability $\frac{1}{2}$ to FIXED.
3. For each agent $i \in \text{STAT}$, query $v_i([m])$. Let $p_0 = \max_{i\in\text{STAT}} v_i([m])$.
4. Draw p uniformly at random from
$$P = \left\{\frac{p_0}{32m^2}, \frac{p_0}{16m^2}, \dots, \frac{p_0}{2}, p_0, 2p_0, \dots, 8m^2 p_0, 16m^2 p_0\right\}.$$
Run a fixed-price auction for agents in FIXED with price p, give any purchased item to the agent who purchased it, collect the corresponding payments, and terminate.

It is easy to check that the above mechanism is universally truthful. If a grand-bundle second price auction happens, truthfulness follows from that of second price auctions. Otherwise, for an agent in STAT, since she will not receive any item anyway, there is no incentive to lie. For an agent i in FIXED, when being asked, her dominant strategy is to purchase her demand set (i.e. a set S that maximizes $v_i(S) - p \cdot |S|$) according to her actual valuation. We prove in the following sections that:

Theorem 2 (Main Theorem). *The generic mechanism is universally truthful, makes exactly one value or demand query to each agent, and returns a* $O(\min(d \log^2 m, \sqrt{m \log m}))$*-approximately optimal allocation of all items in expectation, where* $d = \max_{i\in[n]} \text{SAW}(v_i)$*. When agents have submodular or XOS valuations, the approximation ratio improves to* $O(\log m)$*.*

It may appear that a tighter analysis should give a bound of $\widetilde{O}(\sqrt{d})$, which becomes $\widetilde{O}(1)$ for complement-free agents (when $d = 0$) and $\widetilde{O}(\sqrt{m})$ for general monotone agents (when $d = m - 1$). However, we show that the above bound is in fact almost tight for our protocol, or any protocol within the same framework. Namely,

Proposition 1. *There exist* $2m/(d+1)$ *agents with* SAW-d *valuations such that the generic mechanism yields a* $\Omega(\min\{d, m/d\})$*-approximately optimal allocation.*

We postpone the proof of the above proposition to Appendix B.

3 Warmup: Constraint Homogeneous Valuations

As a warmup, we first prove an approximation guarantee of the generic framework when agents are interested in only disjoint bundles of items. The proof will also be the backbone of the limited-complementarity and general valuation cases to be discussed later. Formally, we are interested in agents with the following class of valuations:

Definition 3 (d-Constraint Homogeneous Valuations [18]). *A valuation* v *is* d*-constraint homogeneous (*d*-CH) if there exists a value* p *(the price-per-item), and disjoint sets* $Q_1, \ldots, Q_\ell$*, each of size at most* d*, so that* $v(Q_k) = p\cdot|Q_k|$ *for every* Q_k*, and the value of every set* $S \subseteq [m]$ *is the sum of values of contained* Q_i*'s, i.e.,*

$$v(S) = \sum_{Q_k \subseteq S} v(Q_k) = p \sum_{Q_k \subseteq S} |Q_k| = p \cdot |\{j : \exists k \text{ s.t. } j \in Q_k \subseteq S\}|.$$

We prove that the generic mechanism gives a $O(d \log m)$ approximation of the optimal welfare when agents have d-CH valuations. We proceed by two cases: when there is an agent whose share in the optimal allocation is large, and when there is no such agent. The former case is directly handled by the grand-bundle second price auction, while the second case requires more effort. All missing proofs in this section are postponed to Appendix C.

Notation. Let $\mathrm{OPT} = (\mathrm{OPT}_1, \ldots, \mathrm{OPT}_n)$ be an optimal allocation, where OPT_i is the set of items that agent i receives. Let $v(\mathrm{OPT}) = \sum_i v_i(\mathrm{OPT}_i)$ be the optimal welfare.

3.1 The Easy Case: When Heavy Agents Exist

First note that:

Lemma 2. *For any* $t \in [0, 1]$*, if for some agent* i*,* $v_i(\mathrm{OPT}_i) \geq \frac{v(\mathrm{OPT})}{t}$*, then a grand-bundle second price auction guarantees welfare at least* $\frac{v(\mathrm{OPT})}{t}$*.*

Therefore, if there is an agent i whose share in the optimal allocation is at least $v_i(\mathrm{OPT}_i) \geq \frac{v(\mathrm{OPT})}{\log m}$, with probability $\frac{1}{2}$ a grand-bundle second price auctions happens, in which case the welfare is at least $\frac{v(\mathrm{OPT})}{\log m}$. The expected welfare is hence at least $\frac{v(\mathrm{OPT})}{2 \log m}$.

3.2 A Thought Experiment: Posted Prices Given Complete Information

Before proceeding to the hard case, we first consider a scenario where the valuations of all agents are known. We demonstrate that in such a case, there exist prices, using which a posted-price auction achieves a d-approximation of the optimal welfare when agents have d-CH valuations. The result does not directly imply a welfare guarantee of our mechanism. Nevertheless, the argument is instrumental for later discussion. We also note that the result in this subsection for the complete information case is not a novel contribution of this paper: for example, a similar statement appears in [11]. We present the entire argument here mainly to provide intuition about the hard case and to be self-contained.

Posted-price auctions. A posted price auction is similar to a fixed price auction, except that the prices for different items can be different. A price is assigned to each item before the auction begins. During the auction, agents are approached in some arbitrary order. Upon being asked, each agent can purchase any subset of the items available, and pay the total prices assigned to these items.

We claim that:

Proposition 2. *For agents with d-CH valuations, there exists prices $\{q_j\}_j$, such that a posted-price auction with prices $\{q_j\}_j$ yields an allocation with welfare at least $\frac{v(\mathrm{OPT})}{2d}$.*

3.3 The Hard Case: When No Heavy Agents Exist

Now we focus on the case where no agent has a share larger than $\frac{v(\mathrm{OPT})}{\log m}$. In such a case, we completely ignore the contribution to the welfare by the second price auction, and analyze solely the contribution of the fixed price auction.

Ideally we would like to run the posted-price auction discussed in the preceding subsection. However, there are two obstacles preventing us from implementing the auction: (1) the valuations of agents are unknown, and (2) computing an optimal allocation is computationally prohibiting. The latter issue can be solved in some sense, by running an approximation algorithm (e.g. [15]), presumably compromising the approximation ratio. On the other hand, there seems to be no easy way around the first issue.

To overcome these difficulties, instead of posting the prices constructed in Proposition 2, our mechanism (1) estimates the interval in which the posted-prices lie, by querying agents in STAT, (2) guesses an appropriate price for agents in FIXED from the estimated interval, and (3) runs a fixed-price auction for agents in FIXED with the price guessed. We show that the expected welfare resulting from such a procedure is not too much worse than the posted-price outcome.

The first step is to show that with high probability, the optimal welfare is relatively equally distributed into STAT and FIXED, so (1) a good approximation restricted to agents in FIXED is also a good approximation with all agents, and (2) an estimation from STAT is useful for guessing the price for FIXED.

Let OPT^{STAT} and $\text{OPT}^{\text{FIXED}}$ be optimal allocations projected to agents in STAT and FIXED respectively. That is, $\text{OPT}_i^{\text{STAT}}$ (resp. $\text{OPT}_i^{\text{FIXED}}$) is OPT_i if i belongs to STAT (resp. FIXED), and $\emptyset$ otherwise.

Lemma 3. *If for some $t \geq 1$, for all $i \in [n]$, $v_i(\text{OPT}_i) \leq \frac{v(\text{OPT})}{t}$, then with probability $1 - 2e^{-t/8}$, $v(\text{OPT}^{\text{STAT}}) \geq \frac{v(\text{OPT})}{4}$ and $v(\text{OPT}^{\text{FIXED}}) \geq \frac{v(\text{OPT})}{4}$.*

Corollary 1. *If for all $i \in [n]$, $v_i(\text{OPT}_i) \leq \frac{v(\text{OPT})}{\log m}$, then with probability $1 - O(1/m)$, $v(\text{OPT}^{\text{STAT}}) \geq \frac{v(\text{OPT})}{4}$ and $v(\text{OPT}^{\text{FIXED}}) \geq \frac{v(\text{OPT})}{4}$.*

We now condition everything on the event (denoted by $\mathcal{E}$) that (1) with probability 1/2, agents are divided into 2 groups, and (2) with probability $1 - O(1/m)$, the two groups are roughly balanced. We only need to show, that when $\mathcal{E}$ happens, the expected welfare of the mechanism is $\Omega\left(\frac{v(\text{OPT})}{d \log m}\right)$.

Let OPT' be an allocation obtained by removing from $\text{OPT}^{\text{FIXED}}$ any item allocated to an agent whose price-per-item is no larger than $\frac{v(\text{OPT}^{\text{FIXED}})}{2m}$. Observe that

Lemma 4. $v(\text{OPT}') \geq \frac{1}{2} v(\text{OPT}^{\text{FIXED}})$.

This means we can safely ignore agents with low price-per-item without losing too much.

For prices high enough, the next lemma shows that we can estimate and guess them with relatively high probability.

Lemma 5. *Conditioned on $\mathcal{E}$, for any $m \geq 512$, price $q \in \left[\frac{v(\text{OPT}^{\text{FIXED}})}{2m^2}, 4v(\text{OPT}^{\text{FIXED}})\right]$, with probability $\frac{1}{|P|} \geq \frac{1}{5 \log m}$, the price p guessed in step 4 of the mechanism satisfies $\frac{1}{4} q \leq p < \frac{1}{2} q$.*

The next step is to show that the fixed-price auction approximates the sum of values of agents whose price-per-item is close to the guessed price p.

Lemma 6. *Conditioned on $\mathcal{E}$, the welfare of the allocation given by the fixed-price auction with price p is at least*

$$\frac{1}{4d} \sum_{i \in \text{FIXED}, \frac{1}{4} p_i \leq p \leq \frac{1}{2} p_i} v_i(\text{OPT}'_i).$$

We are ready to prove a lower bound on the expected welfare of the fixed-price auction.

Lemma 7. *Conditioned on $\mathcal{E}$, the expected welfare generated by the fixed-price auction is $\Omega\left(\frac{v(\text{OPT})}{d \log m}\right)$.*

Now we can put everything together and conclude:

Proposition 3. *When agents have d-CH valuations, the generic mechanism guarantees $O(d \log m)$-approximation of the optimal welfare.*

Proof. When there is a heavy agent (i.e., an agent i with $v_i(\mathrm{OPT}_i) \geq \frac{v(\mathrm{OPT})}{\log m}$), Lemma 2 guarantees expected welfare $\frac{v(\mathrm{OPT})}{2 \log m}$. When there is no heavy agent, Corollary 1 and Lemma 7 guarantee expected welfare $\Omega\left(\frac{v(\mathrm{OPT})}{d \log m}\right)$.

4 Valuations with Limited Complementarity

We show in this section, that for general valuations, the approximation guarantee of the generic mechanism degrades smoothly as the degree of complementarity grows. To establish this result, we first show that if a class of valuations $\mathcal{V}$ is approximated by disjoint bundle valuations with limited bundle size, then the mechanism gives a reasonable guarantee with valuations in $\mathcal{V}$. Then we apply various existing approximation lemmas to establish approximation guarantees of the generic mechanism for submodular, XOS, subadditive, and SAW-d valuations.

Formally, we define pointwise approximation between classes of valuations:

Definition 4 (Pointwise Approximation [4]). *A valuation class $\mathcal{V}$ is pointwise β-approximated by a valuation class $\mathcal{V}'$ if for any valuation $v \in \mathcal{V}$ and for any set $S \subseteq [m]$, there exists a valuation $v' \in \mathcal{V}'$ such that $\beta \cdot v'(S) \geq v(S)$ and for all $T \subseteq [m]$ it holds that $v'(T) \leq v(T)$. We also say such a v' β-approximates v at S.*

We first show that if d-CH β-approximates $\mathcal{V}$, then the generic mechanism guarantees $O(\beta d \log m)$-approximation of the optimal welfare, and then extend the result to d-DB valuations, a superclass of d-CH, as defined below.

Definition 5 (d-Disjoint Bundle Valuations). *A valuation v is d-disjoint bundle (d-DB) if there exists disjoints sets of size at most d and corresponding values $(Q_1, v(Q_1)), \ldots, (Q_\ell, v(Q_\ell))$, so that the value of every set $S \subseteq [m]$ is the sum of values of contained Q_i's, i.e.,*

$$v(S) = \sum_{Q_k \subseteq S} v(Q_k).$$

We first prove the d-CH version of the extension lemma, which plays a central part in our argument:

Lemma 8. *When agents have valuations in class $\mathcal{V}$, for $\beta \leq m$, if $\mathcal{V}$ is pointwise β-approximated by d-CH valuations, then the generic mechanism guarantees $O(\beta d \log m)$-approximation of the optimal welfare.*

Proof. The overall plan is similar to the one discussed in Section 3, except that when no heavy agent exists, we instead use the welfare under the d-CH valuations that approximate the actual valuations as the benchmark, losing a factor of β:

Outline of the proof.

1. If there is a heavy agent whose share in the optimal welfare is at least $\frac{v(\mathrm{OPT})}{\log m}$, the grand-bundle second price auction gives a good approximation.
2. Otherwise, with probability $1/2 - O(1/m)$ event $\mathcal{E}$ happens: the mechanism proceeds to the fixed-price auction, and the optimal welfare is distributed roughly equally to STAT and FIXED. That is, $v(\mathrm{OPT}^{\mathrm{STAT}}) \geq \frac{v(\mathrm{OPT})}{4}$ and $v(\mathrm{OPT}^{\mathrm{FIXED}}) \geq \frac{v(\mathrm{OPT})}{4}$.
3. We construct OPT', the benchmark, as follows:
   (a) For every $i \in \mathrm{FIXED}$, let v_i' be a d-CH valuation, satisfying (1) for all $S \subseteq [m]$, $v_i'(S) = v_i'(S \cap \mathrm{OPT}_i^{\mathrm{FIXED}})$, (2) for all $S \subseteq [m]$, $v_i'(S) \leq v_i(S)$, and (3) $v_i'(\mathrm{OPT}_i^{\mathrm{FIXED}}) \geq \frac{v_i(\mathrm{OPT}_i^{\mathrm{FIXED}})}{\beta}$. In other words, v_i' β-approximates v_i at $\mathrm{OPT}_i^{\mathrm{FIXED}}$. Let p_i be the price-per-item of valuation v_i'. Note that

$$v'(\mathrm{OPT}^{\mathrm{FIXED}}) = \sum_{i \in \mathrm{FIXED}} v_i'(\mathrm{OPT}_i^{\mathrm{FIXED}}) \geq \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{\beta}.$$

   (b) Remove all agents i whose price-per-item p_i is small. That is,

$$\mathrm{OPT}_i' = \begin{cases} \mathrm{OPT}_i^{\mathrm{FIXED}}, & p_i \geq \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2\beta m} \\ \emptyset, & \text{otherwise} \end{cases}.$$

   Note that

$$v'(\mathrm{OPT}') = \sum_{i \in \mathrm{FIXED}} v_i'(\mathrm{OPT}_i') \geq \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2\beta}.$$

4. Let $\mathrm{FIXED}^p = \left\{ i \in \mathrm{FIXED} \mid \frac{1}{4} p_i \leq p < \frac{1}{2} p_i \right\}$. We show that a fixed-price auction with price p generates welfare at least

$$\frac{1}{4d} \sum_{i \in \mathrm{FIXED}^p} v_i'(\mathrm{OPT}_i').$$

5. For $\beta \leq m$, for every $q \in \left[\frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2\beta m}, v(\mathrm{OPT}^{\mathrm{FIXED}}) \right]$, there is some $p' \in P$, which is guessed with probability $\frac{1}{5 \log m}$, such that $\frac{1}{4} q \leq p' \leq \frac{1}{2} q$. Taking the expectation gives the desired approximation ratio. That is,

$$\frac{1}{5 \log m} \cdot \frac{1}{4d \log m} v'(\mathrm{OPT}') \geq \frac{1}{40 \beta d \log m} v(\mathrm{OPT}^{\mathrm{FIXED}}) \geq \frac{1}{160 \beta d \log m} v(\mathrm{OPT}).$$

Consider the outline above. Lemma 2 justifies Step 1. Corollary 1 justifies Step 2. The fact that d-CH pointwise approximates $\mathcal{V}$ justifies Step 3(a). An argument similar to the proof of Lemma 4 justifies Step 3(b). Lemma 5 and an argument similar to the proof of Lemma 7 justify Step 5. We only need to prove the validity of Step 4, which is a relaxed version of Lemma 6.

The proof for Step 4 is again similar to those of Proposition 2 and Lemma 6. For each $i \in \text{FIXED}^p$, we divide items in bundles (induced by v'_i) contained in OPT'_i into two sets:

$$\text{SOLD}_i = \{j \mid \exists k,\, j \in Q^i_k \subseteq \text{OPT}'_i,\, j \text{ is sold}\},$$
$$\text{UNSOLD}_i = \{j \mid \exists k,\, j \in Q^i_k \subseteq \text{OPT}'_i,\, j \text{ is not sold}\}.$$

Note that for $i \notin \text{FIXED}^p$, $\text{SOLD}_i = \text{UNSOLD}_i = \emptyset$. Similarly we define $v'_i(\text{SOLD}_i)$, $v'_i(\text{UNSOLD}_i)$, $v'(\text{SOLD})$ and $v'(\text{UNSOLD})$. We show that the total payment is at least $\frac{v'(\text{SOLD})}{4d}$ and the total buyer surplus is at least $\frac{v'(\text{UNSOLD})}{2}$.

Consider the sold items first. If some item in a bundle Q^i_k is sold, then some agent has to pay $p \geq \frac{1}{4}p_i \geq \frac{v'_i(Q^i_k)}{4d}$ for the item. Summing over all sold bundles, we see that the total payment is at least

$$\sum_{i \in \text{FIXED}^p, Q^i_k \subseteq \text{SOLD}_i} \frac{v'_i(Q^i_k)}{4d} = \sum_i \frac{v'_i(\text{SOLD}_i)}{4d} = \frac{v'(\text{SOLD})}{4d}.$$

Now consider the unsold items. Recall that all items in UNSOLD_i are available throughout the auction. In particular, they are available when agent $i \in \text{FIXED}^p$ chooses the items to purchase. By purchasing exactly the set UNSOLD_i, agent i has a surplus of

$$\begin{aligned}
& v_i(\text{UNSOLD}_i) - p|\text{UNSOLD}_i| \\
\geq\ & v'_i(\text{UNSOLD}_i) - p|\text{UNSOLD}_i| \geq v'_i(\text{UNSOLD}_i) - \frac{p_i|\text{UNSOLD}_i|}{2} \\
=\ & p_i|\text{UNSOLD}_i| - \frac{p_i|\text{UNSOLD}_i|}{2} = \frac{p_i|\text{UNSOLD}_i|}{2} = \frac{v'_i(\text{UNSOLD}_i)}{2}.
\end{aligned}$$

Now since agent i chooses to buy another set, it must be the case that the set she purchases gives at least the same amount of surplus, i.e., $\frac{v'_i(\text{UNSOLD}_i)}{2}$. Summing over agents, we see that the total surplus is at least $\frac{v'(\text{UNSOLD})}{2}$.

It follows that the welfare is at least

$$\frac{v'(\text{SOLD})}{4d} + \frac{v'(\text{UNSOLD})}{2} \geq \frac{1}{4d} \sum_{i \in \text{FIXED}^p} v'_i(\text{OPT}'_i).$$

This concludes the proof for Step 4 and the theorem.

Now observe that the above argument can be easily modified to work if we replace d-CH valuations with d-DB valuations. Formally,

Lemma 9. *When agents have valuations in class $\mathcal{V}$, for $\beta \leq m$, if $\mathcal{V}$ is pointwise β-approximated by d-DB valuations, then the generic mechanism guarantees $O(\beta d \log m)$-approximation of the optimal welfare.*

We postpone the proof of Lemma 9 to Appendix D. Note that we do not need to know the d-CH or d-DB valuations which approximate the v_i's — the existence of the approximation suffices for our purpose.

With Lemma 9, we are now ready to translate the approximation lemmas by d-DB to welfare guarantees of the generic mechanism. Restricted to complement-free classes, it is well known that:

Lemma 10 (Folklore). *Fractionally subadditive (or XOS) valuations are pointwise* 1*-approximated by* 1*-DB (i.e. additive) valuations.*

Dobzinski [5] and Devanur et al. [4] independently show that:

Lemma 11 ([5, 4]). *Subadditive valuations are pointwise* $O(\log m)$*-approximated by* 1*-CH (i.e. homogeneously additive) valuations.*

And beyond complement-free classes, Chen et al. [3] show that:

Lemma 12 ([3]). *For any* $d \geq 1$*, the class* SAW-d *is pointwise* $2H_m$*-approximated by* $2d$*-CH, where* $H_i = \sum_{k \in [i]} \frac{1}{k}$ *is the* i*-th harmonic number.*

Applying Lemma 9 to Lemmas 10, 11, and 12, we obtain:

Theorem 3. *When agents have (1) submodular or XOS, (2) subadditive, or (3)* SAW-d *valuations, the generic mechanism guarantees (1)* $O(\log m)$*-, (2)* $O(\log^2 m)$*-, or (3)* $O(d \log^2 m)$*-approximation of the optimal welfare, respectively.*

Proof. The Theorem follows from Lemma 9 by setting β to (1) 1, (2) $O(\log m)$, and (3) $2H_m = O(\log m)$, and d to (1) 1, (2) 1, and (3) $2d'$ respectively.

5 General Monotone Valuations

In this section, we show that the generic mechanism guarantees $O(\sqrt{m \log m})$-approximation of the optimal welfare, thereby concluding the proof of Theorem 2. We do this, again, by modifying the outline given in Section 4 (proof deferred to Appendix E).

Theorem 4. *When agents have monotone valuations, the generic mechanism guarantees* $O(\sqrt{m \log m})$*-approximation of the optimal welfare.*

Putting Theorems 3 and 4 together, Theorem 2 follows directly.

References

1. Abraham, I., Babaioff, M., Dughmi, S., Roughgarden, T.: Combinatorial auctions with restricted complements. In: Proceedings of the 13th ACM Conference on Electronic Commerce. pp. 3–16. ACM (2012)
2. Assadi, S., Singla, S.: Improved truthful mechanisms for combinatorial auctions with submodular bidders. In: 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS). pp. 233–248. IEEE (2019)

3. Chen, W., Teng, S.H., Zhang, H.: Capturing complementarity in set functions by going beyond submodularity/subadditivity. In: 10th Innovations in Theoretical Computer Science Conference (ITCS 2019). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2019)
4. Devanur, N., Morgenstern, J., Syrgkanis, V., Weinberg, S.M.: Simple auctions with simple strategies. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation. pp. 305–322. ACM (2015)
5. Dobzinski, S.: Two randomized mechanisms for combinatorial auctions. In: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, pp. 89–103. Springer (2007)
6. Dobzinski, S.: An impossibility result for truthful combinatorial auctions with submodular valuations. In: Proceedings of the forty-third annual ACM symposium on Theory of computing. pp. 139–148. ACM (2011)
7. Dobzinski, S.: Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In: Proceedings of the forty-eighth annual ACM symposium on Theory of Computing. pp. 940–948. ACM (2016)
8. Dobzinski, S., Nisan, N., Schapira, M.: Approximation algorithms for combinatorial auctions with complement-free bidders. In: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing. pp. 610–618. ACM (2005)
9. Dobzinski, S., Nisan, N., Schapira, M.: Truthful randomized mechanisms for combinatorial auctions. In: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing. pp. 644–652. ACM (2006)
10. Dobzinski, S., Vondrák, J.: The computational complexity of truthfulness in combinatorial auctions. In: Proceedings of the 13th ACM Conference on Electronic Commerce. pp. 405–422. ACM (2012)
11. Düetting, P., Feldman, M., Kesselheim, T., Lucier, B.: Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In: Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on. pp. 540–551. IEEE (2017)
12. Eden, A., Feldman, M., Friedler, O., Talgam-Cohen, I., Weinberg, S.M.: A simple and approximately optimal mechanism for a buyer with complements. In: Proceedings of the 2017 ACM Conference on Economics and Computation. pp. 323–323 (2017)
13. Ehsani, S., Hajiaghayi, M., Kesselheim, T., Singla, S.: Prophet secretary for combinatorial auctions and matroids. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 700–714. SIAM (2018)
14. Feige, U.: On maximizing welfare when utility functions are subadditive. SIAM Journal on Computing **39**(1), 122–142 (2009)
15. Feige, U., Feldman, M., Immorlica, N., Izsak, R., Lucier, B., Syrgkanis, V.: A unifying hierarchy of valuations with complements and substitutes. In: Twenty-Ninth AAAI Conference on Artificial Intelligence (2015)
16. Feige, U., Izsak, R.: Welfare maximization and the supermodular degree. In: Proceedings of the 4th conference on Innovations in Theoretical Computer Science. pp. 247–256. ACM (2013)
17. Feige, U., Vondrak, J.: Approximation algorithms for allocation problems: Improving the factor of 1-1/e. In: null. pp. 667–676. IEEE (2006)
18. Feldman, M., Friedler, O., Morgenstern, J., Reiner, G.: Simple mechanisms for agents with complements. In: Proceedings of the 2016 ACM Conference on Economics and Computation. pp. 251–267. ACM (2016)

19. Feldman, M., Gravin, N., Lucier, B.: Combinatorial auctions via posted prices. In: Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms. pp. 123–135. SIAM (2014)
20. Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proceedings of the forty-fourth annual ACM symposium on Theory of computing. pp. 123–136. ACM (2012)
21. Krysta, P., Vöcking, B.: Online mechanism design (randomized rounding on the fly). In: International Colloquium on Automata, Languages, and Programming. pp. 636–647. Springer (2012)
22. Lavi, R., Swamy, C.: Truthful and near-optimal mechanism design via linear programming. Journal of the ACM (JACM) **58**(6), 1–24 (2011)
23. Mirrokni, V., Schapira, M., Vondrák, J.: Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In: Proceedings of the 9th ACM conference on Electronic commerce. pp. 70–77. ACM (2008)
24. Nisan, N.: The communication complexity of approximate set packing and covering. In: International Colloquium on Automata, Languages, and Programming. pp. 868–875. Springer (2002)
25. Syrgkanis, V., Tardos, E.: Composable and efficient mechanisms. In: Proceedings of the forty-fifth annual ACM symposium on Theory of computing. pp. 211–220. ACM (2013)
26. Vondrák, J.: Optimal approximation for the submodular welfare problem in the value oracle model. In: Proceedings of the fortieth annual ACM symposium on Theory of computing. pp. 67–74. ACM (2008)

A Further Related Work

There is an extremely rich body of research on truthful combinatorial auctions, agents with complements, and pointwise approximation. Our results and techniques are closely related to all these fields and the research therein. We give in this subsection a comprehensive overview of results and techniques related to our work.

A.1 Truthful Combinatorial Auctions

Most related to our work are the line of work on computationally efficient truthful combinatorial auctions [9, 5, 7, 2]. They propose and develop the following powerful framework for truthful combinatorial auctions with demand queries:

- If there is an agent whose share in the optimal allocation is large enough, then we can sell the grand bundle via a second-price auction to this agent, or someone with a higher value, to guarantee a good welfare.
- Otherwise we divide all agents into 2 groups. We then query one group to gather information, and sell to the other group by posting prices to items to ensure truthfulness.

In this paper, we simplify this framework, and show that what we believe to be an essential part in fact gives smoothly degrading welfare guarantees as the degree of complementarity grows, for all monotone valuation functions simultaneously.

Lavi and Swamy [22] give an $O(\sqrt{m})$-approximate truthful in expectation protocol for agents with general valuations, and Abraham et al. [1] give a $O(\log^k m)$-approximate truthful in expectation protocol for agents with PH-k valuations[8]. These results are incomparable to ours for the following reasons. First, these protocols are truthful in expectation, as opposed to being universally truthful. For reasons discussed above, such protocols may not be able to prevent agents from misreporting. Second, all PH-k valuations (as considered in [1]) are superadditive, which means even if we are willing to settle with a loss of $\Omega(\log^m m)$, their protocol still works only for a highly restricted subclass of all monotone valuations, i.e., the superadditive class.

Feldman et al. [19] consider a Bayesian setting, where agents' valuations are drawn independently from publicly known distributions. They give a k-approximation posted-price protocol w.r.t. any allocation algorithm, for agents with MPH-k valuations[9]. Plugging in the k-approximation algorithm in [15], their result implies a $O(k^2)$-approximation protocol for MPH-k agents. The $O(k^2)$ bound was later improved by Dütting et al. [11] to $O(k)$, as a corollary of a more general framework introduced therein. The protocols in both work proceed by querying the prior distributions, posting prices on items, and then letting agents buy their demand bundles. They are therefore trivially truthful. Notably, the above two papers use similar revenue-utility decomposition techniques as used in this paper to derive welfare bounds[10]. However, these results are incomparable to ours, since they rely crucially on access to the prior distributions of valuations, and the approximation guarantees are w.r.t. the expected optimal welfare given the prior distributions. In contrast, our protocol provides welfare guarantees for all possible valuations in the respective classes.

A.2 Valuations with Complements

Valuations with complements have been considered in various settings. Most related to our results is the paper by Chen et al. [3], where they introduce and justify the SAW hierarchy and its counterpart, the SMW hierarchy which generalizes submodular valuations. In particular, they observe that different problems often require different parametrizations by different measures of complementarity, and the SAW hierarchy in particular is useful in analyzing the PoA of simple auctions. Their findings suggest that the SAW hierarchy could also be useful in analyzing other dynamics with strategic agents, e.g., truthful combinatorial auctions, which we confirm in this paper. Abraham et al. [1], Feige et al. [15] and Chen et al. [3] consider separately the algorithmic problem of maximizing welfare among agents with limited complements. Different measures of complementarity

[8] PH-k valuations are valuations represented as nonnegatively weighted hypergraphs with hyperedges of size not exceeding k, where the value of a set is the sum of the weights of all hyperedges contained in this set.

[9] A valuation is MPH-k [15] iff it is the pointwise maximum of some PH-k valuations.

[10] Similar arguments also appear in a number of other papers on related problems. See, e.g., [20, 13].

have been used to analyze the PoA [15, 18, 3] and revenue [12] of simple auctions. It is worth noting that Devanur et al. [4] define the Constraint Homogeneous (CH) class, which was later generalize to d-CH by Feldman et al. [18], to facilitate the analysis of PoA of simple auctions. The d-CH class is also an important auxiliary valuation class in [3], and this paper.

A.3 Pointwise Approximation

Devanur et al. [4] formally define the notion of pointwise approximation, and show that it preserves smoothness as defined by Syrgkanis and Tardos [25], and therefore PoA bounds from the smoothness framework. Built upon this extension lemma, pointwise approximation has been used extensively to study the PoA of simple auctions [4, 15, 18, 3]. These PoA bounds are established essentially in the same way: prove smoothness for a simple class of valuations, establish pointwise approximation of more complex classes by the simple class, and apply the extension lemma due to Devanur et al. [4]. Little was known about the properties of pointwise approximation outside the smoothness framework. We note that our Lemma 9 is not simply a translation of the extension lemma in [4]. In particular, Lemma 9 explicitly requires approximation by d-DB valuations (as opposed to arbitrary valuations in [4]) since the proof would not work otherwise.

B Omitted Proofs in Section 2

Proof (Proof of Proposition 1). W.l.o.g. suppose $m = (d+1) \cdot \ell$ for some $\ell \in \mathbb{N}$. We divide all items into ℓ groups $M_1, \ldots M_\ell$, each of size $d+1$. For group k, we create 2 agents, $2k-1$ and $2k$, who are interested only in items in M_k. Agent $2k-1$ is unit-demand, and has value 1 for each item in M_k. Agent $2k$ has value d if he has the entire M_k, and 0 otherwise. The optimal welfare is $\ell \cdot d = \frac{dm}{d+1}$, achieved by allocating M_k to agent $2k$. However, if we sell the entire bundle to a single agent, then the maximum possible welfare is d. If we run the fixed-price auction, then let p be the posted price. If $p < 1$, then for each k, agent $2k-1$ buys an arbitrary item in M_k, resulting in a welfare of ℓ. If $p \geq 1$, then no agent buys anything throughout the auction, resulting in a welfare of 0. In any case, the welfare produced by the mechanism is no larger than $\max\{d, \ell\}$, and the approximation ratio is $\Omega(\min\{d, m/d\})$.

C Omitted Proofs in Section 3

Proof (Proof of Lemma 2). Let i_1 be the agent such that $v_{i_1}(\mathrm{OPT}_{i_1}) \geq \frac{v(\mathrm{OPT})}{t}$. By truthfulness of second price auctions, some agent, say i_2, with the maximum value for the grand bundle, wins all items. By monotonicity, we have

$$v_{i_2}([m]) \geq v_{i_1}([m]) \geq v_{i_1}(\mathrm{OPT}_{i_1}) \geq \frac{v(\mathrm{OPT})}{t},$$

as desired.

Proof (Proof of Proposition 2). We first construct the prices used in the posted-price auction. Fix an optimal allocation $\{\mathrm{OPT}_i\}_i$. Let agent i's valuation be given by the price-per-item p_i and bundles $\{Q^i_1, \ldots, Q^i_{\ell_i}\}$. For each item $j \in \mathrm{OPT}_i$, if the bundle containing j is allocated to agent i as a whole (i.e. there is some bundle Q^i_k, such that $j \in Q^i_k \subseteq \mathrm{OPT}_i$), let the price of j in the posted-price auction be $q_j := \frac{p_i}{2}$. Otherwise, let $q_j := 0$.

We now show that the posted-price auction yields a reasonable approximation of the optimal welfare. Consider all bundles completely allocated to an agent i in OPT. For each Q^i_k of these bundles, one of the two happens: either some item $j \in Q^i_k$ is purchased by some agent (not necessarily by i), or all items in the bundle remain unsold at the end of the auction. For each agent i, we partition all items in bundles owned by agent i in OPT into two sets: SOLD_i and UNSOLD_i, containing items in the two kinds of bundles respectively. Formally,

$$\mathrm{SOLD}_i = \{j \mid \exists k,\ j \in Q^i_k \subseteq \mathrm{OPT}_i,\ j \text{ is sold}\},$$
$$\mathrm{UNSOLD}_i = \{j \mid \exists k,\ j \in Q^i_k \subseteq \mathrm{OPT}_i,\ j \text{ is not sold}\}.$$

Let $v(\mathrm{SOLD}) = \sum_i v_i(\mathrm{SOLD}_i)$, and $v(\mathrm{UNSOLD}) = \sum_i v_i(\mathrm{UNSOLD}_i)$. Note that $v_i(\mathrm{OPT}_i) = v_i(\mathrm{SOLD}_i) + v_i(\mathrm{UNSOLD}_i)$, and $v(\mathrm{OPT}) = v(\mathrm{SOLD}) + v(\mathrm{UNSOLD})$.

Consider the sold items first. We show that the total payment made by the agents is at least $\frac{v(\mathrm{SOLD})}{2d}$. If some item in a bundle Q^i_k is sold, then some agent has to pay $\frac{p_i}{2} \geq \frac{v_i(Q^i_k)}{2d}$ for the item. Recall that all bundles are disjoint, and the valuation over them is therefore in a sense additive. Summing over all sold bundles, we see that the total price paid by all agents during the auction is at least

$$\sum_{i\in[n], Q^i_k \subseteq \mathrm{SOLD}_i} \frac{v_i(Q^i_k)}{2d} = \sum_i \frac{v_i(\mathrm{SOLD}_i)}{2d} = \frac{v(\mathrm{SOLD})}{2d}.$$

Now consider the unsold items. We show that in the allocation given by the posted-price auction, the surplus (i.e., the total value of the items a buyer gets minus the total price she pays) of agent i is at least $\frac{v_i(\mathrm{UNSOLD}_i)}{2}$. Note that all items in UNSOLD_i are available throughout the auction. In particular, they are available when agent i chooses the items to purchase. By purchasing exactly the set UNSOLD_i, agent i has a surplus of

$$v_i(\mathrm{UNSOLD}_i) - \frac{p_i|\mathrm{UNSOLD}_i|}{2} = p_i|\mathrm{UNSOLD}_i| - \frac{p_i|\mathrm{UNSOLD}_i|}{2} = \frac{p_i|\mathrm{UNSOLD}_i|}{2} = \frac{v_i(\mathrm{UNSOLD}_i)}{2}.$$

Now since agent i chooses to buy another set, it must be the case that the set she purchases gives at least the same amount of surplus, i.e., $\frac{v_i(\mathrm{UNSOLD}_i)}{2}$. Summing over agents, we see that the total surplus is at least $\frac{v(\mathrm{UNSOLD})}{2}$.

Now note that the welfare of all agents is the sum of the total payment and the total surplus, which is at least $\frac{v(\mathrm{SOLD})}{2d} + \frac{v(\mathrm{UNSOLD})}{2} \geq \frac{v(\mathrm{OPT})}{2d}$.

Proof (Proof of Lemma 3). We show that $v(\mathrm{OPT}^{\mathrm{STAT}}) = \sum_{i \in \mathrm{STAT}} v_i(\mathrm{OPT}_i) \geq \frac{v(\mathrm{OPT})}{t}$ with probability at least $1 - e^{-t/8}$. The case of FIXED is totally symmetric, and the lemma follows from a union bound. Let $X_i := \mathbb{I}[i \in \mathrm{STAT}]$. By definition of the mechanism $\{X_i\}_i$ are i.i.d.

Consider concentration of $\sum_{i\in[n]} X_i \cdot v_i(\mathrm{OPT}_i)$. Note that

– Each summand $X_i \cdot v_i(\mathrm{OPT}_i) \in [0, v_i(\mathrm{OPT}_i)]$.
– All summands are independent.
– $\mathbb{E}\left[\sum_{i\in[n]} X_i \cdot v_i(\mathrm{OPT}_i)\right] = \frac{1}{2}v(\mathrm{OPT})$.

By Hoeffding bound,

$$\Pr\left[\frac{1}{2}v(\mathrm{OPT}) - \sum_{i\in[n]} X_i \cdot v_i(\mathrm{OPT}_i) \le \frac{1}{4}v(\mathrm{OPT})\right] \le \exp\left(-\frac{2(v(\mathrm{OPT})/4)^2}{\sum_{i\in[n]} v_i(\mathrm{OPT}_i)^2}\right)$$
$$\le \exp\left(-\frac{v(\mathrm{OPT})^2/8}{t\left(\frac{v(\mathrm{OPT})}{t}\right)^2}\right) = \exp\left(-\frac{t}{8}\right),$$

which concludes the proof.

Proof (Proof of Lemma 4).

$$v(\mathrm{OPT}^{\mathrm{FIXED}}) - v(\mathrm{OPT}') = \sum_{i:p_i \le \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m}} v_i(\mathrm{OPT}_i^{\mathrm{FIXED}}) \le \sum_{i:p_i \le \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m}} p_i|\mathrm{OPT}_i^{\mathrm{FIXED}}|$$
$$\le \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m} \sum_{i:p_i \le \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m}} |\mathrm{OPT}_i^{\mathrm{FIXED}}| \le \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m} \cdot m$$
$$= \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2},$$

as desired.

Proof (Proof of Lemma 5). Recall the procedure by which p is chosen: first query $p_0 = \max_{i\in\mathrm{STAT}} v_i([m])$, and then choose p uniformly at random from $P = \left\{\frac{p_0}{32m^2}, \frac{p_0}{16m^2}, \dots, 8m^2p_0, 16m^2p_0\right\}$. First note that once we know p_0, there are at most $5\log m$ possible values of p, and each value is chosen with probability $\frac{1}{5\log m}$. We only need to show there is some price $p' \in P$, such that $\frac{1}{4}q \le p' \le \frac{1}{2}q$. Or, strengthening the condition, we want to show that $q \in \left[\frac{p_0}{8m^2}, 16m^2p_0\right]$. Observe that conditioned on $\mathcal{E}$ and given Lemma 4,

$$p_0 = \max_{i\in\mathrm{STAT}} v_i([m]) \ge \frac{1}{m}\sum_{i\in\mathrm{STAT}} v_i(\mathrm{OPT}_i^{\mathrm{STAT}}) \ge \frac{1}{4m}v(\mathrm{OPT}) \ge \frac{1}{4m}v(\mathrm{OPT}^{\mathrm{FIXED}}) \ge \frac{1}{16m}q.$$

On the other hand,

$$q \ge \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m^2} \ge \frac{v(\mathrm{OPT})}{8m^2} \ge \frac{v(\mathrm{OPT}^{\mathrm{STAT}})}{8m^2} \ge \frac{p_0}{8m^2}.$$

We conclude that $q \in \left[\frac{p_0}{8m^2}, 16mp_0\right] \subseteq \left[\frac{p_0}{8m^2}, 16m^2p_0\right]$. The lemma follows.

Proof (Proof of Lemma 6). The proof is overall similar to that of Proposition 2. Let FIXED^p be the set of agents in FIXED, whose price-per-item is close to p. That is,

$$\text{FIXED}^p = \left\{ i \in \text{FIXED} \mid \frac{1}{4}p_i \leq p < \frac{1}{2}p_i \right\}.$$

For each $i \in \text{FIXED}^p$, we divide items in bundles contained in OPT'_i into two sets:

$$\text{SOLD}_i = \{j \mid \exists k,\, j \in Q^i_k \subseteq \text{OPT}'_i,\, j \text{ is sold}\},$$
$$\text{UNSOLD}_i = \{j \mid \exists k,\, j \in Q^i_k \subseteq \text{OPT}'_i,\, j \text{ is not sold}\}.$$

Note that for $i \notin \text{FIXED}^p$, $\text{SOLD}_i = \text{UNSOLD}_i = \emptyset$. Similarly we define $v_i(\text{SOLD}_i)$, $v_i(\text{UNSOLD}_i)$, $v(\text{SOLD})$ and $v(\text{UNSOLD})$. The goal is again to show that the total payment is at least $\frac{v(\text{SOLD})}{4d}$ and the total buyer surplus is at least $\frac{v(\text{UNSOLD})}{2}$.

Consider the sold items first. If some item in a bundle Q^i_k is sold, then some agent has to pay $p \geq \frac{1}{4}p_i \geq \frac{v_i(Q^i_k)}{4d}$ for the item. Summing over all sold bundles, we see that the total payment is at least

$$\sum_{i \in \text{FIXED}^p, Q^i_k \subseteq \text{SOLD}_i} \frac{v_i(Q^i_k)}{4d} = \sum_i \frac{v_i(\text{SOLD}_i)}{4d} = \frac{v(\text{SOLD})}{4d}.$$

Now consider the unsold items. Recall that all items in UNSOLD_i are available throughout the auction. In particular, they are available when agent $i \in \text{FIXED}^p$ chooses the items to purchase. By purchasing exactly the set UNSOLD_i, agent i has a surplus of

$$v_i(\text{UNSOLD}_i) - p|\text{UNSOLD}_i| \geq v_i(\text{UNSOLD}_i) - \frac{p_i|\text{UNSOLD}_i|}{2} = p_i|\text{UNSOLD}_i| - \frac{p_i|\text{UNSOLD}_i|}{2}$$
$$= \frac{p_i|\text{UNSOLD}_i|}{2} = \frac{v_i(\text{UNSOLD}_i)}{2}.$$

Now since agent i chooses to buy another set, it must be the case that the set she purchases gives at least the same amount of surplus, i.e., $\frac{v_i(\text{UNSOLD}_i)}{2}$. Summing over agents, we see that the total surplus is at least $\frac{v(\text{UNSOLD})}{2}$.

It follows that the welfare is at least

$$\frac{v(\text{SOLD})}{4d} + \frac{v(\text{UNSOLD})}{2} \geq \frac{1}{4d} \sum_{i \in \text{FIXED}^p} v_i(\text{OPT}'_i),$$

which concludes the proof.

Proof (Proo of Lemma 7). Note that:

- For any $p_1, p_2 \in P = \left\{\frac{p_0}{32m^2}, \frac{p_0}{16m^2}, \ldots, 8m^2p_0, 16m^2p_0\right\}$ where $p_1 \neq p_2$, $\text{FIXED}^{p_1} \cap \text{FIXED}^{p_2} = \emptyset$.

– Letting $\text{FIXED}' = \cup_{p' \in P} \text{FIXED}^{p'}$, $\sum_{i \in \text{FIXED}'} v_i(\text{OPT}'_i) = v(\text{OPT}')$. This is because for any $i \in \text{FIXED}$ where $\text{OPT}'_i \neq \emptyset$, $\frac{v(\text{OPT}^{\text{FIXED}})}{2m} \leq p_i \leq v(\text{OPT}^{\text{FIXED}})$. p_i therefore falls into the interval covered by Lemma 5.

The expected welfare is therefore at least

$$\sum_{p' \in P} \Pr[p = p'] \cdot \frac{1}{4d} \sum_{i \in \text{FIXED}^{p'}} v_i(\text{OPT}'_i) = \frac{1}{|P|} \frac{v(\text{OPT}')}{4d} \geq \frac{v(\text{OPT}')}{20d \log m} \geq \frac{v(\text{OPT})}{160d \log m}.$$

The first equality and the first inequality follow from Lemma 5 and the two observations above. The second inequality follows from Lemmas 3 and 4.

D Omitted Proofs in Section 4

Proof (Proof of Lemma 9). We modify Step 3 of the outline, and argue that the other steps still work.

The new Step 3. For every $i \in \text{FIXED}$, let $w_i = \{(Q^i_k, w_i(Q^i_k)\}_{k \in [\ell_i]}$ be the d-DB valuation that approximates v_i at $\text{OPT}^{\text{FIXED}}_i$, satisfying (1) for all $S \subseteq [m]$, $w_i(S) = w_i(S \cap \text{OPT}^{\text{FIXED}}_i)$, (2) for all $S \subseteq [m]$, $w_i(S) \leq v_i(S)$, and (3) $w_i(\text{OPT}^{\text{FIXED}}_i) \geq \frac{v_i(\text{OPT}^{\text{FIXED}}_i)}{\beta}$. For each $(Q^i_k, w_i(Q^i_k))$, construct a dummy single-minded agent a^i_k who has value $w_i(Q^i_k)$ for bundle Q^i_k and is not interested in anything else. That is, the valuation of a^i_k satisfy

$$v'_{a^i_k}(S) = \mathbb{I}[Q^i_k \subseteq S] \cdot w_i(Q^i_k).$$

Note that for any $i \in \text{FIXED}$, we have:

1. For all $S \subseteq [m]$, $w_i(S) = \sum_{k \in [\ell_i]} v'_{a^i_k}(S)$.
2. For all $S \subseteq [m]$, $\sum_{k \in [\ell_i]} v'_{a^i_k}(S) = w_i(S) \leq v_i(S)$.
3. $\sum_{k \in [\ell_i]} v'_{a^i_k}(\text{OPT}^{\text{FIXED}}_i) = w_i(\text{OPT}^{\text{FIXED}}_i) \geq \frac{v_i(\text{OPT}^{\text{FIXED}}_i)}{\beta}$.

Now $v'(\text{OPT}^{\text{FIXED}})$ is undefined. We instead move on to constructing OPT' with agents $\{a^i_k\}_{k,i}$. Note that $v'_{a^i_k}$ is a d-CH valuation. Let $p_{a^i_k} = \frac{w_i(Q^i_k)}{|Q^i_k|}$ be the price-per-item of dummy agent a^i_k. Let OPT' be such that

$$\text{OPT}'_i = \begin{cases} Q^i_k, p_{a^i_k} \geq \frac{v(\text{OPT}^{\text{FIXED}})}{2\beta m} \\ \emptyset, \quad \text{otherwise} \end{cases}.$$

We claim that with agents $\{a^i_k\}$, valuations $\{v'_{a^i_k}\}$, and benchmark OPT', Steps 4 and 5 still work and yield the desired guarantee. Most parts of the original argument carries over directly, except that we need to check the buyer surplus bound more carefully. Recall that $\text{FIXED}^p = \left\{a^i_k \mid \frac{1}{4}p \leq p_{a^i_k} < \frac{1}{2}p\right\}$. For each $a^i_k \in \text{FIXED}^p$, if some item in Q^i_k is sold, then $\text{UNSOLD}_{a^i_k} = \emptyset$. Otherwise let $\text{UNSOLD}_{a^i_k} = Q^i_k$. Note that for $a \notin \text{FIXED}^p$, $\text{SOLD}_a = \text{UNSOLD}_a = \emptyset$.

For each actual agent i, let A^p_i be the indices of dummy agents whose price-per-item is close to p. That is,

$$A^p_i = \left\{k \mid a^i_k \in \text{FIXED}^p\right\}.$$

Let $\text{UNSOLD}_i = \cup_{k\in A^p_i} \text{UNSOLD}_{a^i_k}$. Agent i can always choose to purchase set UNSOLD_i. Her surplus is therefore at least

$$\begin{aligned}
v_i(\text{UNSOLD}_i) - p|\text{UNSOLD}_i| &\geq w_i(\text{UNSOLD}_i) - p|\text{UNSOLD}_i| \\
&= \sum_{k\in A^p_i} v'_{a^i_k}(\text{UNSOLD}_i) - p|\text{UNSOLD}_i| \\
&= \sum_{k\in A^p_i} \left(v'_{a^i_k}(\text{OPT}'_{a^i_k}) - p|Q^i_k|\right) \\
&\geq \sum_{k\in A^p_i} \left(p_{a^i_k}|Q^i_k| - \frac{1}{2}p_{a^i_k}|Q^i_k|\right) \\
&= \sum_{k\in A^p_i} \frac{1}{2} v'_{a^i_k}(\text{OPT}'_{a^i_k}).
\end{aligned}$$

Note that $\{A^p_i\}_i$ is a partition of FIXED^p. Summing over i gives that the surplus is at least

$$\frac{1}{2}\sum_{i\in\text{FIXED}}\sum_{k\in A^p_i} v'_{a^i_k}(\text{OPT}'_{a^i_k}) = \frac{1}{2}\sum_{a\in\text{FIXED}^p} v'_a(\text{UNSOLD}_a) = \frac{1}{2}v'(\text{UNSOLD}),$$

which concludes the proof.

E Omitted Proofs in Section 5

Proof (Proof of Theorem 4). We give the modified outline first:

Outline of the proof.

1. If there is a heavy agent whose share in the optimal welfare is at least $\frac{v(\text{OPT})}{16\sqrt{m\log m}}$, the grand-bundle second price auction gives a good approximation.
2. Otherwise, with probability $1/2 - O(1/e^m)$ event $\mathcal{E}$ happens: the mechanism proceeds to the fixed-price auction, and the optimal welfare is distributed roughly equally to STAT and FIXED. That is, $v(\text{OPT}^{\text{STAT}}) \geq \frac{v(\text{OPT})}{4}$ and $v(\text{OPT}^{\text{FIXED}}) \geq \frac{v(\text{OPT})}{4}$.

3. (a) We construct OPT', the benchmark, as follows: Let $p_i = \frac{v(\mathrm{OPT}_i^{\mathrm{FIXED}})}{|\mathrm{OPT}_i^{\mathrm{FIXED}}|}$ if $\mathrm{OPT}_i^{\mathrm{FIXED}} \neq \emptyset$, and 0 otherwise. For $i \in \mathrm{FIXED}$,

$$\mathrm{OPT}'_i = \begin{cases} \mathrm{OPT}_i^{\mathrm{FIXED}}, & p_i \geq \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m} \text{ and } |\mathrm{OPT}_i^{\mathrm{FIXED}}| \leq \sqrt{\frac{m}{\log m}} \\ \emptyset, & \text{otherwise} \end{cases}.$$

   (b) For $i \in \mathrm{FIXED}$, let v'_i be such that

$$v'_i(S) = \begin{cases} \mathbb{I}[\mathrm{OPT}_i^{\mathrm{FIXED}} \subseteq S] \cdot v_i(\mathrm{OPT}_i^{\mathrm{FIXED}}), & \mathrm{OPT}'_i \neq \emptyset \\ 0, & \text{otherwise} \end{cases}.$$

   Note that for any $i \in \mathrm{FIXED}$, v'_i is $\sqrt{\frac{m}{\log m}}$-CH, and v'_i 1-approximates v_i at OPT'_i. We show later that $v'(\mathrm{OPT}') \geq \frac{v(\mathrm{OPT})}{16}$.
4. Let $\mathrm{FIXED}^p = \left\{ i \in \mathrm{FIXED} \mid \frac{1}{4} p_i \leq p < \frac{1}{2} p_i \right\}$. A fixed-price auction with price p generates welfare at least

$$\frac{\sqrt{\log m}}{4\sqrt{m}} \sum_{i \in \mathrm{FIXED}^p} v'_i(\mathrm{OPT}'_i).$$

5. For $\beta \leq m$, for every $q \in \left[\frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2m}, v(\mathrm{OPT}^{\mathrm{FIXED}}) \right]$, there is some $p' \in P$, which is guessed with probability $\frac{1}{5 \log m}$, such that $q \in \mathrm{FIXED}^{p'}$. Taking the expectation gives the desired approximation ratio. That is,

$$\frac{1}{5 \log m} \cdot \frac{\sqrt{\log m}}{4\sqrt{m}} v'(\mathrm{OPT}') \geq \frac{v(\mathrm{OPT})}{320\sqrt{m \log m}}.$$

Consider the outline above. Lemma 2 justifies Step 1. By plugging $t = \sqrt{\frac{m}{\log m}}$ in, Lemma 3 justifies Step 2. Steps 4 and 5 are totally similar to the corresponding parts in the proof of Lemma 8. We only need to show that in Step 3, conditioned on $\mathcal{E}$, the benchmark constructed approximates the optimal welfare. That is, $v'(\mathrm{OPT}') \geq \frac{v(\mathrm{OPT})}{16}$.

Observe that in OPT', we eliminate the shares of two kinds of agents: (1) agents whose price-per-item is too low, and (2) agents who have too large shares in terms of cardinality. We bound the two parts of the loss separately, and show that a significant fraction of the optimal welfare remains. First note that a similar argument to the proof of Lemma 4 establishes that the first part of the loss is at most $\frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2}$. For the second part, note that the number of agents whose shares are large is at most $\frac{m}{\sqrt{m/\log m}} = \sqrt{m \log m}$. Since there is no heavy agent, the total value of the shares of these agents is at most

$$\sqrt{m \log m} \cdot \frac{v(\mathrm{OPT})}{16\sqrt{m \log m}} = \frac{v(\mathrm{OPT})}{16}.$$

Thus we have

$$v'(\mathrm{OPT}') \geq v(\mathrm{OPT}^{\mathrm{FIXED}}) - \frac{v(\mathrm{OPT}^{\mathrm{FIXED}})}{2} - \frac{v(\mathrm{OPT})}{16} \geq \frac{v(\mathrm{OPT})}{8} - \frac{v(\mathrm{OPT})}{16} = \frac{v(\mathrm{OPT})}{16},$$

where the second inequality follows from the definition of event $\mathcal{E}$. This concludes the proof.