# Combinatorial Auctions via Posted Prices

Michal Feldman[*] Nick Gravin[†] Brendan Lucier[‡]

**Abstract**

We study mechanisms for the submodular combinatorial auction problem in a Bayesian setting. In this problem, $m$ indivisible goods are to be allocated among $n$ buyers, whose valuations are drawn independently from known distributions over submodular functions (and, more generally, fractionally subadditive functions). We design a posted-price mechanism that is dominant strategy incentive compatible, runs in time polynomial in $n$ and $m$, and achieves a constant fraction of the expected optimal social welfare.

Our mechanism is straightforward: item prices are determined in advance and do not change, then the consumers approach the seller sequentially in an arbitrary order, each purchasing any of her favorite bundles from among the unsold items at the posted prices. Given black-box access to any allocation algorithm $A$ and sample access to the prior distribution, we show how to compute prices that yield at least half of the expected welfare of $A$. Finally, our results extend to valuations with complements, where the approximation factor degrades linearly with the level of complementarity.

[*]Tel-Aviv University; mfeldman@tau.ac.il
[†]Microsoft Research; ngravin@microsoft.com
[‡]Microsoft Research; brlucier@microsoft.com

# 1 Introduction

The canonical problem in market design is to efficiently allocate a set of resources among a set of self-interested agents. Such allocation problems range in scope from the trade of a single item between a seller and a buyer, to combinatorial auctions in which many heterogeneous goods are to be divided among multiple participants with complex and idiosyncratic preferences. Scenarios of the latter type have attracted significant recent attention from the computer science community, due to algorithmic challenges presented by the underlying allocation problem. For example, the efficient allocation of cloud resources involves the scheduling of computing tasks, and the allocation of wireless spectrum involves finding large independent sets in graphs that represent interference constraints. The primary challenge in algorithmic mechanism design is to marry algorithmic solutions to such problems with the economic principles that underpin market design.

This tension between algorithmic and economic constraints is exemplified by the well-studied submodular combinatorial auction (CA) problem. The goal in this problem is to efficiently partition a collection of indivisible goods among a group of consumers with submodular preferences. This problem admits strong solutions from both the algorithmic and economic perspectives: a simple greedy algorithm obtains a worst-case constant approximation, and the VCG mechanism implements the optimal outcome and is dominant strategy incentive compatible (DSIC). However, it remains a vexing open problem to design a DSIC mechanism with sub-polynomial worst-case approximation factor.[1]

In addition to computational efficiency, the submodular CA problem imposes two strong requirements: implementation in dominant strategies, and good performance in the worst case over realized consumer preferences. In this work we relax the latter of these, the worst-case evaluation criterion. We adopt instead a Bayesian approach to welfare maximization. We assume that consumer preferences are drawn independently from distributions known to the mechanism designer, whose goal is to maximize welfare in expectation over the type realizations. Notably, while the mechanism can be prior-dependent, we still require that its incentive properties hold ex post.

For this setting, we design a DSIC mechanism that generates a constant fraction of the optimal expected welfare, for arbitrary distributions over submodular valuations (in fact, for the more general class of fractionally subadditive valuations). Prior to this work, no DSIC mechanism was known to give constant approximation to the Bayesian submodular CA problem. Moreover, our mechanism takes the simple and natural form of *posted prices*. The items for sale are initially assigned prices (that never change throughout the process). The participants then sequentially consume their favorite bundles.

We proceed with the formal model and results, followed by a discussion of our mechanism's properties and implementation.

## 1.1 Our model

Our setting consists of a set $M$ of $m$ indivisible objects and a set of $n$ buyers. Each buyer has a valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$ that indicates his value for every set of objects. We assume valuations are monotone non-decreasing, normalized so that $v_i(\emptyset) = 0$, and scaled to lie in $[0, 1]$.[2] The profile of buyer valuations is denoted by $\mathbf{v} = (v_1, \ldots, v_n)$. An *allocation* of $M$ is a vector of sets $\mathbf{X} = (X_1, \ldots, X_n)$, where $X_i$ denotes the bundle assigned to buyer $i \in [n]$, and $X_i \cap X_k = \emptyset$

[1] It is known that no DSIC mechanism can achieve a sub-polynomial worst-case approximation under the value query model of communication [17], but the problem remains open for demand queries.

[2] The assumption that valuations are bounded is used only to bound sampling error when computing prices, and is unnecessary if sampling is unnecessary (i.e., if valuation distributions are given explicitly).

for every $i \neq k$ (note that it is not required that all items are allocated). The *social welfare* of an allocation $\mathbf{X}$ is $\mathsf{SW}(\mathbf{X}, \mathbf{v}) = \sum_{i=1}^{n} v_i(X_i)$, and the optimal welfare is denoted by $\mathrm{OPT}(\mathbf{v})$. When clear from context we omit $\mathbf{v}$ and write $\mathsf{SW}$ and OPT for the social welfare and optimal welfare, respectively.

Fix an implicit valuation profile $\mathbf{v}$. Suppose items are associated with prices $p_1, \ldots, p_m$. The utility of buyer $i$ being allocated bundle $X_i$ under item prices $\mathbf{p} = (p_1, \ldots, p_m)$ is $u_i(X_i, \mathbf{p}) = v_i(X_i) - \sum_{j \in X_i} p_j$. Given prices $\mathbf{p}$, the *demand correspondence* $D_i(M, \mathbf{p})$ of buyer $i$ contains the sets of objects that maximize buyer $i$'s utility; i.e., $D_i(M, \mathbf{p}) = \operatorname{argmax}_{S \subseteq M} u_i(S, \mathbf{p})$.

We consider a *Bayesian* setting, where the bidders' valuations are drawn independently from distributions $\mathcal{F}_1, \ldots, \mathcal{F}_n$. Write $\mathcal{F} = \mathcal{F}_1 \times \cdots \times \mathcal{F}_n$, so that $\mathbf{v}$ is drawn from $\mathcal{F}$. We think of $\mathcal{F}$ as being public knowledge, whereas the realization $v_i$ is known only to agent $i$. In the Bayesian framework, an allocation $\mathbf{X}$ is said to be an $\alpha$-approximation (for social welfare) if

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{F}}\left[\mathsf{SW}(\mathbf{X}, \mathbf{v})\right] \geq (1/\alpha) \cdot \mathbf{E}_{\mathbf{v} \sim \mathcal{F}}\left[\mathrm{OPT}(\mathbf{v})\right].$$

### 1.1.1 Valuation Classes

We study both complement-free valuations, and valuations that exhibit complementarities.

There is a standard hierarchy of complement-free valuations (see [27]): additive $\subset$ gross substitutes $\subset$ submodular $\subset$ XOS $\subset$ subadditive.

***Additive:*** there exist values $v(\{1\}), \ldots, v(\{m\})$ such that $v(S) = \sum_{j \in S} v(\{j\})$ for all $S \subset M$.

***Gross substitutes* (GS):** the demand for any item does not decrease if the price of other items increase. Formally, for any item $j \in S$, where $S \in D(M, \mathbf{p})$ and for any price vector $\mathbf{q} \succ \mathbf{p}$ such that $q_j = p_j$ there should be $T \in D(M, \mathbf{q})$ such that $j \in T$.

***Submodular*:** for every $S \subseteq T \subseteq M$ and $j \in M$, $v(j|T) \leq v(j|S)$, where $v(j|S) \coloneqq v(S \cup \{j\}) - v(S)$.

***XOS*:** there exists a collection of additive functions $a_1(\cdot), \ldots, a_k(\cdot)$ such that for every set $S \subseteq M$, $v(S) = \max_{1 \leq i \leq k} a_i(S)$.[3]

To study valuations with complements, we consider the hierarchy *maximum over positive hypergraphs* **(MPH)**, introduced recently by [19]. This hierarchy is general enough to encapsulate all monotone valuation functions, and its level captures the degree of complementarity. We defer a formal description to Section 3.

### 1.1.2 Computational models

An algorithm for the combinatorial auction problem receives as input a valuation profile $\mathbf{v}$, and returns an allocation profile. We write $\mathcal{A}$ for an algorithm, and $\mathcal{A}(\mathbf{v})$ for the allocation returned. As any explicit description of $v_i : 2^M \to \mathbb{R}_{\geq 0}$ would have size exponential in $m$, it is usually assumed that there is an oracle access to $v_i$. We consider the following oracles:

- **Value oracle** takes as input a set $T$, and returns $v_i(T)$;
- **Demand oracle** takes as an input a price vector $\mathbf{p}$, and returns a set from the demand correspondence $D_i(M, \mathbf{p})$, breaking ties arbitrarily but consistently;
- **XOS oracle** (only for XOS function $v_i$) takes as input a set $T$, and returns the corresponding additive representative function for the set $T$, i.e., an additive function $a_i(\cdot)$ such that (i) $v_i(S) \geq a_i(S)$ for any $S \subset [m]$, and (ii) $v_i(T) = a_i(T)$;

[3] The class XOS is also referred to as *fractionally subadditive* valuations [18].

Most of our results require access to demand oracles. Our perspective is that the demand oracle captures the most basic decision problem a buyer faces in a market with item prices. Since a primary focus of this paper is pricing mechanisms, we assume throughout that consumers are able to address this decision problem, and take demand oracle access for granted.

## 1.2 Our results

The posted price mechanisms considered in this paper proceed in the following steps:

1. Construct a price vector $\mathbf{p}$, independently of the buyers' reports.
2. To each buyer $i$, in an arbitrary order, allocate his preferred bundle from among the items not already allocated. Each buyer pays the sum of item prices for the bundle he receives.

Prices are chosen independently of the buyers' reported valuations, and do not change as buyers are served. Since each agent's declaration is used only to select the set to allocate, and since a utility-maximizing set is always chosen, we can immediately conclude that the mechanism is truthful and individually rational (see Theorem 2.1). The challenge is to set prices that guarantee a good approximation to the optimal welfare. The strength of the approximation guarantee we establish comes from two requirements: we require that the approximation holds (a) for *any* order of agent arrival (i.e., the mechanism is *order-oblivious* in the sense of [12]), and (b) for *any* way agents break ties among bundles in their demands. Buyers arrive in an arbitrary order, each purchases a bundle in her demand among the remaining items, breaking ties in an arbitrary manner. Our results will hold for any choice of buyer order, and regardless of the manner in which ties are broken.

Our main result is that posted price mechanism can achieve a constant approximation to the optimal welfare in polynomial time, for distributions over fractionally subadditive valuations. In fact, given an arbitrary allocation algorithm $\mathcal{A}$, one can design a mechanism that yields nearly half of the welfare of $\mathcal{A}$.

**Main Theorem: [2-approximation for XOS]** Fix any $\varepsilon > 0$ and suppose we have black-box access to an allocation algorithm $\mathcal{A}$, sample access to $\mathcal{XOS}$ distributions $\mathcal{F}$, and $\mathcal{XOS}$ and demand oracles for the valuations in the support of $\mathcal{F}$. Then there is a posted price mechanism with expected social welfare at least $\frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathcal{A}(\mathbf{v})] - \varepsilon$ that runs in time $\text{POLY}(n, m, 1/\varepsilon)$.

We establish our main theorem by showing how to compute an appropriate vector of prices, given requisite access to the type distributions. We then show that using these prices in a posted-price mechanism yields the claimed expected welfare, regardless of the buyer order and regardless of how ties are broken. The factor of 2 is tight, even for a single item, and even if the mechanism can choose the arrival order and the tie-breaking method; see Section 2.3.

The main theorem is stated as a reduction: given an allocation algorithm that achieves acceptable welfare for the type distributions at hand, one can generate a posted price mechanism that performs nearly as well. One can apply the reduction to known worst-case approximation algorithms, to obtain constant-factor approximation mechanisms for arbitrary distributions of $\mathcal{XOS}$ valuations. For some subclasses of valuations, the approximation factor can be improved by using tailored algorithms, or the oracle requirements can be relaxed using known reductions. Table 1 lists some notable applications.

| | | Oracle Model | |
|---|---|---|---|
| **Valuation Class** | **Approximation** | **Computing Prices** | **Full Mechanism** |
| $\mathcal{XOS}$ | $\frac{1}{2}(1-1/e)$ | $\mathcal{XOS}$ & Demand Oracles | $\mathcal{XOS}$ & Demand Oracles |
| Submodular | $\frac{1}{2}(1-1/e)$ | Value Oracle | Demand Oracle |
| Gross Substitutes | $\frac{1}{2}$ | Value Oracle | Value Oracle |

Table 1: Applications of the Main Theorem. Approximation factors achievable for different valuation classes, by applying known allocation algorithms [18, 30, 7] and oracle model equivalences [9, 29] for each class.

The following features of our result deserve particular mention:

- **Strong IC guarantees:** (i) DSIC: Even though our mechanism is prior-dependent, it is incentive compatible ex post (i.e., in dominant strategies). This is a stronger notion than Bayesian incentive compatibility (BIC), which requires only that truth telling is a Bayes-Nash equilibrium. In particular, the incentive properties hold regardless of the beliefs held by the agents: truthtelling is utility-optimal even if the agents have a more refined signal of buyer preferences than the mechanism. (ii) Group strategyproof: Our mechanism is *weakly group strategyproof*, meaning that no coalition of agents can deviate in a way that strictly gains each one of them.
- **(Truly) poly-time:** The running time of our mechanism is independent of the space of agent types. In particular, the running time is polynomial in $n$ and $m$, even if each agent's type distribution is supported on exponentially many valuations. This is important because the valuation space can plausibly be exponentially large in $m$.[4]

**An Illustrating Example.** To give some insight into our results, consider the case of a single item and $n$ bidders with values drawn i.i.d. from some distribution $F$. In this setting, the Vickrey auction generates the efficient outcome. How well can one approximate the efficient outcome by setting a single price $p$ (that depends only on $F$ and $n$) and allocating to a random bidder[5] with value greater than $p$, if any exists?

As it turns out, there is a simple pricing scheme for a single item that yields half of the optimal social welfare. In fact, this is a direct application of the Prophet inequality; see [25]. One solution, described also by Kleinberg and Weinberg [26], is to set price $p$ equal to half of the expected highest value. To see why this works, write $\omega$ for the probability that the item is sold, and write $i^*$ for the agent with the highest value for the item (note $i^*$ is a random variable). The expected welfare of the auction is simply the expected revenue plus the expected buyer surplus (that is, the utility of the winner). The expected revenue from the auction is precisely $p \cdot \omega = \frac{1}{2}\mathbb{E}[v_{i^*}] \cdot \omega$. On the other hand, since the item is unsold with probability $(1-\omega)$, buyer $i^*$ always has at least this probability of seeing the item. The expected buyer surplus is therefore at least $\mathbb{E}[(1-\omega)(v_{i^*}-p)] = \frac{1}{2}\mathbb{E}[v_{i^*}] \cdot (1-\omega)$. Putting this together, the expected welfare of the auction is at least $\frac{1}{2}\mathbb{E}[v_{i^*}] \cdot \omega + \frac{1}{2}\mathbb{E}[v_{i^*}] \cdot (1-\omega) = \frac{1}{2}\mathbb{E}[v_{i^*}]$, as claimed.

[4] Consider, for example, a base valuation modified by independent noise on the value of each item.

[5] Since agents are iid, this is equivalent to an adversarial order that doesn't depend on the value realizations

Our main result essentially shows that the reasoning above for the single-item example can be extended to markets with multiple heterogeneous items and asymmetric buyers, as long as buyer preferences lie in the class of $\mathcal{XOS}$ valuations. In this sense, our result can be interpreted as a multiple-item extension of the prophet inequality.

So far we have only considered complement-free valuations. Our results extend to more general functions, with approximation factor and runtime bound that degrades with the level of complementarity in the valuations. We express our result in terms of MPH-$k$ valuations [19] — a hierarchy that spans all monotone valuations, parameterized by the complementarity level $k$. A formal definition of MPH-$k$ appears in Section 3.

**Theorem:** [**$k$-approximation for MPH−k**] Fix any $\varepsilon > 0$, and suppose we have black-box access to an allocation algorithm $\mathcal{A}$, sample access to MPH−k distributions $\mathcal{F}$, and MPH−k and demand oracles for the valuations in the support of $\mathcal{F}$. Then there exists a posted price mechanism with expected social welfare at least $\frac{1}{k}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathcal{A}(\mathbf{v})] - \varepsilon$ that runs in time $\text{POLY}(n, m^k, 1/\varepsilon)$.

By the fact that the class of MPH−m valuations (i.e., $k = m$) is equivalent to the class of all monotone functions [19], the last theorem implies the existence of posted prices that gives expected social welfare of at least $\frac{1}{\Omega(m)} \cdot \text{OPT}$. In Appendix B we show that this bound is tight.

### 1.3 Discussion: Implementation and Simplicity

Posted-price mechanisms are inherently robust, as they are oblivious to the buyer arrival order and to the way in which ties are broken between utility-maximizing sets. This robustness is suggestive of a natural indirect implementation of posted-price mechanisms. Specifically, one can imagine that a designer sets the item prices, but then consumers arrive in sequence and each simply purchases a desired bundle of goods. Such an implementation has numerous advantages over its direct-revelation analog: it is inherently decentralized, and it enables the buyers to reveal less about their preferences than in a direct revelation mechanism. It is simple and robust: even if the designer is incorrectly informed about the prior distribution, the agents will still behave as price-takers; the choice of prices only impacts the welfare generated. The only burden imposed upon the buyers is that they respond to demand queries, which are regarded natural in many contexts. Finally, it is *obviously strategyproof*, as defined by Li [28], since each agent is simply asked to select their favorite outcome from a menu of possibilities.[6]

The fact that ties can be broken arbitrarily is especially notable when contrasted with common notions of market equilibrium, such as Walrasian equilibrium. In a Walrasian (or competitive) equilibrium, there is a choice of prices such that all agents can *simultaneously* obtain a most-demanded set of goods. This simultaneity is a strong property, and is one that our posted-price mechanisms will not satisfy in general. On the other hand, competitive equilibria rely crucially on ties between different utility-maximizing sets, and it is vital that such ties are broken in the correct way. We therefore view our mechanism as being relevant even when a Walrasian equilibrium can be found, such as when the type distribution is a point mass on a profile of gross substitutes valuations. In such cases, the prices we construct might generate less welfare than competitive equilibrium (which can implement the optimal outcome), but does so more robustly since it does not rely on correct tie-breaking.

[6] Li [28] defines a mechanism to be obviously strategyproofness if it admits dominant strategies, and moreover the worst possible outcome under the dominant strategy is no worse than the best possible outcome under any other action. The indirect implementation of our mechanism satisfies this property (since an agent's outcome is completely determined by her choice), but a direct revelation implementation might not.

A different example of a simple interaction protocol for combinatorial markets is the simultaneous item auctions. In these auctions, items are sold via independent first or second price auctions and each buyer may participate simultaneously in as many auctions as she deems necessary. Such auctions have been adopted by certain Internet marketplace platforms. In recent years a significant body of work has been devoted to the analysis of simultaneous auctions in the Bayesian settings [6, 14, 20, 24]. Neither of first- or second-price auction formats is DSIC, so their performance is typically analyzed under the assumption that the participants bid at equilibrium. Such analysis hinges on the predictiveness of the equilibrium solution concept. While one might expect equilibrium play to be predictive in some circumstances, it is known that individual participants in a game do not always play at equilibrium, for a variety of reasons. For example, finding a Bayesian Nash Equilibrium can be computationally intractable in simultaneous item auctions [10]. In contrast, our construction illustrates that if one permits prior-dependence, then this difficulty can be overcome by designing explicit prices and relieving buyers from the task of bidding strategically.

### 1.4 Related work

Our work relates to the design of truthful submodular combinatorial auctions. It is known that no sub-polynomial factor approximation is possible under the value query model [17], or for succinctly-described valuations [16]. Given access to demand queries, a randomized truthful $O(\log m \log \log m)$ worst-case approximation exists [15]. It is a major open question whether there is a truthful constant-factor mechanism, using demand queries. We show that in Bayesian settings, where the performance of the mechanism is evaluated based on its expected social welfare given an input distribution, one can indeed design a truthful constant-factor submodular combinatorial auction.

Ours is not the first work to turn to the Bayesian setting for combinatorial auction design. Hartline and Lucier [23], Hartline et al. [22] and Bei and Huang [4] provide black-box reductions that convert an arbitrary welfare-maximization algorithm into an (approximately) Bayesian incentive compatible (BIC) mechanism, without loss of welfare. However, these results relax the dominant strategy IC requirement to Bayesian Incentive compatibility (BIC). Also, these reductions require time polynomial in the support size of an agent's valuation distribution, which can be exponential in $n$ and $m$. In contrast, our mechanism runs in time polynomial in $n$ and $m$, regardless of the valuation distributions' support sizes. Alaei [2] presents a general method for designing DSIC combinatorial auction mechanisms in Bayesian settings, using an algorithm for a related single-agent optimization problem, but does not consider submodular CAs.[7]

Another approach in the study of combinatorial auctions calls to abandon incentive compatibility and analyze the efficiency of a mechanism using the Bayesian price of anarchy (BPOA) measure. That is, assume that the agents reach some Bayesian Nash equilibrium (BNE) and compare the welfare achieved at the worst BNE to the optimal welfare. Particularly relevant to us is the line of work on Bayesian combinatorial auctions analyzing the BPOA of *simultaneous item* auctions, see, e.g., [6, 14, 20, 24]. A constant BPOA has been shown, even when valuations are subadditive [20]. Our posted prices mechanism also achieves constant approximation guarantees for $\mathcal{XOS}$ valuations, but differs in that it is dominant strategy incentive compatible and prior-dependent.

There is a long line of research studying the performance of posted price mechanisms under the objective of maximizing revenue. When there is only a single item for sale, posted prices obtain 78% of the optimal revenue in large markets [8]. When agents have unit-demand preferences, a

[7] While Alaei [2] does not explicitly discuss submodular CAs, our understanding is that one could use his methodology together with an algorithm for (single-agent) submodular function maximization to construct a constant-factor DSIC mechanism for the submodular CA problem. We note that such a mechanism would not fall within the posted-price paradigm.

form of posted-price mechanism extracts a constant fraction of the optimal revenue [11, 12, 13]. When there is only a single item for sale, this constant factor persists even when the distributions are unknown, as long as they satisfy the monotone hazard rate assumption [3].

# 2 Posted Prices for XOS Valuations

This section is dedicated to the construction and analysis of our posted-price mechanism for $\mathcal{XOS}$ valuations. We begin by describing the general form of the mechanism and formally establishing its properties, independent of the choice of prices. We will then show how to compute prices that guarantee high social welfare in expectation.

## 2.1 Posted Price Mechanisms

A direct revelation implementation of the posted price mechanism is listed as Algorithm 1. This implementation supposes the existence of item prices, chosen in advance of the execution of the mechanism. In particular, the prices are independent of the declared buyer types. In the mechanism listing, variable $R$ maintains the set of unallocated items, initially $M$. Recall that $D_i(\mathbf{p}, R)$ is the set of utility-maximizing subsets of $R$ for agent $i$, given item prices $\mathbf{p}$.

**ALGORITHM 1:** Posted-Price Mechanism, parametrized by price vector $\mathbf{p}$.

Set $R = M$.
Choose an arbitrary permutation $\pi$ of the bidders.
For each buyer $i = \pi(1), \ldots, \pi(n)$:
  Choose $Y_i$ to be an arbitrary set from $D_i(\mathbf{p}, R)$.
  Update $R \leftarrow R \backslash Y_i$.
Return allocation $\mathbf{Y}$. Buyer $i$ pays $\sum_{j \in Y_i} p_j$.

Since each agent's declaration is used only to select the set to allocate, and since a utility-maximizing set is always chosen, we can immediately conclude that the mechanism is truthful and individually rational. In fact, it must also be weakly group strategyproof, in the sense that no coalition of agents can jointly deviate to benefit each one of the coalition members. After all, the first member of any coalition (in order $\pi$) who deviates from choosing a utility-maximizing set cannot get a higher utility, by definition. Finally, the mechanism makes a linear number of demand queries. The following theorem summarizes these properties.

**Theorem 2.1.** *The posted price mechanism is dominant strategy incentive compatible, weakly group strategyproof, and ex post individually rational. The algorithm makes $n$ demand queries, and runs in polynomial time given access to a demand oracle.*

## 2.2 Computing Prices

The heart of our mechanism is the computation of appropriate prior-dependent prices that guarantee high welfare in expectation. The following theorem is the main result of this section.

**Theorem 2.2.** *Let distribution $\mathcal{F}$ over $\mathcal{XOS}$ valuation profiles be given via a sample access to $\mathcal{F}$. Suppose that we have*

1. *black-box access to an allocation algorithm $\mathcal{A}$, and*

2. *an $\mathcal{XOS}$ query oracle (for valuations sampled from $\mathcal{F}$).*

*Then, for any $\varepsilon > 0$, we can compute item prices in time* $\text{POLY}(m, n, 1/\varepsilon)$ *such that, for any buyer arrival order, the expected welfare of the posted price mechanism is at least* $\frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \varepsilon$.

**Implications.** Before proving Theorem 2.2, let us discuss its implications, which are summarized in Table 1. First, note that using an $\alpha$-approximation algorithm for $\mathcal{A}$ in Theorem 2.2 results in a posted price mechanism with approximation factor $\alpha/2$, minus an additive error term that can be made as small as desired.

If we assume access to demand oracles, then we can use the poly-time algorithm of Feige [18] with approximation factor $1 - 1/e$ as a black box. Theorem 2.2 combined with Theorem 2.1 implies the existence of a DSIC mechanism with expected social welfare at least $\frac{1}{2}(1 - 1/e)\text{OPT} - \varepsilon$ and runtime $\text{POLY}(m, n, 1/\varepsilon)$.

For submodular valuations, one could instead use the algorithm by Vondrak [30] with tight approximation factor $\alpha = 1 - 1/e$ that utilizes only value queries. Since $\mathcal{XOS}$ queries can be simulated by value queries for submodular valuations [9], we obtain the following corollary:

**Corollary 2.3.** *Given sample access to submodular distributions $\mathcal{F}$ and value oracle access to each valuation in the support of $\mathcal{F}$, for every $\varepsilon > 0$, one can compute item prices in time* $\text{POLY}(m, n, 1/\varepsilon)$*, such that, for any buyer arrival order, the expected welfare of the posted price mechanism is at least* $\frac{1}{2}(1 - 1/e)\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \varepsilon$.

As before, one can implement the mechanism from Corollary 2.3 as a direct revelation mechanism, if one also has access to demand oracles for the valuations.

For gross substitutes valuations, demand queries can be implemented with a polynomial number of value queries [29], and an optimal allocation can be computed in polynomial time using demand queries [7]. The following corollary follows:

**Corollary 2.4.** *Given sample access to gross substitutes distributions $\mathcal{F}$ and value oracle access to each valuation in the support of $\mathcal{F}$, for every $\varepsilon > 0$, one can compute item prices in time* $\text{POLY}(m, n, 1/\varepsilon)$*, such that, for any buyer arrival order, the expected welfare of the posted price mechanism is at least* $\frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \varepsilon$.

Here, the mechanism from Corollary 2.3 can be implemented as a direct revelation mechanism, using only value queries.

### 2.2.1 Warmup: Deterministic Valuations.

Before proving Theorem 2.2 in general, it will be useful to consider the special case where $\mathcal{F}$ is a point-mass distribution and $\mathcal{A}$ is deterministic. The following claim is a restatement of Theorem 2.2, restricted to this special case.

**Claim 2.5.** *Given any $\mathcal{XOS}$ valuation profile $\mathbf{v}$ and allocation $\mathbf{X}$, there exist item prices such that, for any buyer arrival order, the welfare of the posted price mechanism is at least $\frac{1}{2}SW(\mathbf{X})$. Morover, such prices can be found in polynomial time, given $\mathcal{XOS}$ oracle access to the valuations.*

*Proof.* We can assume without loss of generality that $\mathbf{X}$ allocates all items. Moreover, since the valuations are $\mathcal{XOS}$, there exists for each $i$ an additive valuation $a_i$ such that $a_i(S) \leq v_i(S)$ for all $S$, and $a_i(X_i) = v_i(X_i)$. We will use these additive valuations to define the item prices: for each $i$ and each $j \in X_i$, we will choose the price of item $j$ to be $p_j = \frac{1}{2}a_i(\{j\})$. Note that these prices are easily computable given the allocation $\mathbf{X}$ and $\mathcal{XOS}$ oracles for the agent valuations.

We now observe two things about this choice of prices. First,

$$\sum_{j \in M} p_j = \frac{1}{2} SW(\mathbf{X}). \tag{1}$$

We can therefore interpret the price of item $j$ as half of that item's contribution of the social welfare of $\mathbf{X}$. Second, for any agent $i$, and any subset $Q \subseteq X_i$, we have

$$v_i(X_i \backslash Q) + 2 \sum_{j \in Q} p_j \geq v_i(X_i). \tag{2}$$

We can therefore think of the prices as partially covering any loss in value due to some items in $X_i$ being unavailable for agent $i$ to purchase.

Now consider the outcome of the posted price mechanism. Choose some bidder permutation $\pi$, and let $\mathbf{Y}$ be any allocation consistent with the posted price mechanism, given $\pi$ and $\mathbf{p}$. Write $\texttt{SOLD} = \cup_i Y_i$ for the set of items that are sold. Let $\texttt{SOLD}_i = X_i \cap \left(\cup_{\pi(k) < \pi(i)} Y_k\right)$ be the set of items from $X_i$ that are sold before agent $i$ is considered by the mechanism.

We are ready to bound the welfare of the posted price mechanism. Recall that the total welfare is equal to the mechanism's revenue, plus the sum of the agents' utilities. Since agent $i$ has the option of purchasing set $X_i \backslash \texttt{SOLD}_i$, he must obtain at least as much utility as that option, and hence

$$\begin{aligned} u_i(Y_i, \mathbf{p}) &\geq v_i(X_i \backslash \texttt{SOLD}_i) - \sum_{j \in X_i \backslash \texttt{SOLD}_i} p_j \\ &\geq v_i(X_i) - 2 \sum_{j \in \texttt{SOLD}_i} p_j - \sum_{j \in X_i \backslash \texttt{SOLD}_i} p_j \qquad \text{(by (2))} \\ &= \left( v_i(X_i) - \sum_{j \in X_i} p_j \right) - \sum_{j \in \texttt{SOLD}_i} p_j. \end{aligned}$$

Summing over all agents, we get

$$\begin{aligned} \sum_i u_i(Y_i, \mathbf{p}) &\geq \left( \sum_i v_i(X_i) - \sum_j p_j \right) - \sum_i \sum_{j \in \texttt{SOLD}_i} p_j \\ &= \frac{1}{2} \sum_i v_i(X_i) - \sum_i \sum_{j \in \texttt{SOLD}_i} p_j \qquad \text{(by (1))} \\ &\geq \frac{1}{2} \sum_i v_i(X_i) - \sum_{j \in \texttt{SOLD}} p_j \end{aligned}$$

Where the last line follows because each item in $\texttt{SOLD}$ appears in at most one set $\texttt{SOLD}_i$. But now, since the revenue generated by the mechanism is precisely $\sum_{j \in \texttt{SOLD}} p_j$, we conclude that its welfare is

$$\sum_i u_i(Y_i, \mathbf{p}) + \sum_{j \in \texttt{SOLD}} p_j \geq \frac{1}{2} \sum_i v_i(X_i) = \frac{1}{2} SW(\mathbf{X})$$

as required. □

### 2.2.2 The General Case: Proof of Theorem 2.2.

We now proceed with the proof of Theorem 2.2. The proof will proceed in two parts. We begin with Lemma 2.6, which establishes the *existence* of prices that achieve the desired welfare properties, without regard for computation. In fact, Lemma 2.6 establishes something stronger: if the prices are perturbed slightly, this does not have too large an effect on expected welfare. We will then use this stronger property to show how the prices can be computed efficiently via sampling. This sampling process generates the additional additive error term in Theorem 2.2.

Before delving into the details of the proof, we need the following definition of an item's welfare contribution. Fix a valuation profile $\mathbf{v} = (v_1, \ldots, v_n)$ and algorithm $\mathcal{A}$, and let $\mathbf{X} = (X_1, \ldots, X_n)$ be the allocation $\mathcal{A}(\mathbf{v})$. For each $\mathcal{XOS}$ valuation function $v_i$, define the *corresponding additive representative* function for the set $X_i$ as the function $a_i$ satisfying: (i) $v_i(S) \geq a_i(S)$ for any $S \subset [m]$, and (ii) $v_i(X_i) = a_i(X_i)$. For every item $j \in X_i$ we define $\mathsf{SW}_j(\mathbf{v}) := a_i(\{j\})$. We think of $\mathsf{SW}_j(\mathbf{v})$ as the contribution of item $j$ to the social welfare under valuation profile $\mathbf{v}$.

**Lemma 2.6.** *Given a distribution $\mathcal{F}$ over $\mathcal{XOS}$ valuations, let $\mathbf{p}$ be the price vector defined as*

$$p_j = \frac{1}{2} \cdot \mathop{\mathbf{E}}_{\mathbf{v} \sim \mathcal{F}} \left[ SW_j(\mathbf{v}) \right].$$

*Let $\mathbf{p}'$ be any price vector such that $|p'_j - p_j| < \delta$ for all $j$. Then, for any arrival order $\pi$, the expected welfare of Algorithm 1 allocation $\mathbf{Y}$ under prices $\mathbf{p}'$ is at least $\frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[SW(\mathcal{A}(\mathbf{v}))] - m\delta$.*

*Proof.* First, by the definition of $p_j$,

$$p'_j = \mathop{\mathbf{E}}_{\mathbf{v} \sim \mathcal{F}} \left[ \mathsf{SW}_j(\mathbf{v}) - p'_j \right] + 2(p'_j - p_j) = \sum_{i=1}^{n} \mathop{\mathbf{E}}_{\mathbf{v} \sim \mathcal{F}} \left[ \left( \mathsf{SW}_j(\mathbf{v}) - p'_j \right) \cdot \mathbb{1} \left[ j \in X_i(\mathbf{v}) \right] \right] + 2(p'_j - p_j). \quad (3)$$

We are now going to estimate the sum of buyers' utilities in expectation over $\mathcal{F}$. Fix $i$ and $\mathbf{v} = (v_i, \mathbf{v}_{\text{-}i})$. Let $\mathtt{SOLD}_i(\mathbf{v}, \pi)$ denote the set of items that have been sold before the arrival of buyer $i$. Recall that buyer $i$ picks an allocation[8] that maximizes his utility with respect to his valuation $v_i$ and prices $\mathbf{p}$, from among the items in $M \setminus \mathtt{SOLD}_i(\mathbf{v}, \pi)$.

Consider another random valuation profile $\widetilde{\mathbf{v}}_{\text{-}i} \sim \mathcal{F}_{\text{-}i}$ which is independent of $\mathbf{v}$. Let $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ be the allocation returned by $\mathcal{A}$ on input $(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$. We consider additive representative function $a_i$ for the set $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$, so that $a_i(\{j\}) = \mathsf{SW}_j(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ for each $j \in X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$. Let $S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i}) := X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) \setminus \mathtt{SOLD}_i(\mathbf{v}, \pi)$ be the subset of items in $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ that are available to be purchased when buyer $i$ arrives. We note that buyer $i$ could have picked the set $S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})$ and, therefore, his utility must be at least the utility he would get from purchasing that set. Thus we have

$$u_i(\mathbf{v}) \geq \mathop{\mathbf{E}}_{\widetilde{\mathbf{v}}_{\text{-}i}} \left[ \sum_{j \in S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})} \max\left( \mathsf{SW}_j(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) - p'_j, 0 \right) \right].$$

Adding these inequalities for all buyers and taking the expectation over all $\mathbf{v} \sim \mathcal{F}$ we get

$$\mathop{\mathbf{E}}_{\mathbf{v} \sim \mathcal{F}} \left[ \sum_{i=1}^{n} u_i(\mathbf{v}) \right] \geq \sum_{j \in M} \sum_{i=1}^{n} \mathop{\mathbf{E}}_{v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i}} \left[ \max\left( \mathsf{SW}_j(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) - p'_j, 0 \right) \cdot \mathbb{1} \left[ j \in X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) \right] \cdot \mathbb{1} \left[ j \notin \mathtt{SOLD}_i(\mathbf{v}, \pi) \right] \right]. \quad (4)$$

[8] If a buyer has more than one bundle in his demand correspondence, then ties can be broken arbitrarily – even adversarially.

We further observe that $\texttt{SOLD}_i(\mathbf{v},\pi)$ does not depend on $v_i$. That is, $\texttt{SOLD}_i(\mathbf{v},\pi) = \texttt{SOLD}_i(\mathbf{v}_{\text{-}i},\pi)$. Therefore, we can rewrite (4) as follows:

$$\begin{aligned}
\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} u_i(\mathbf{v})\right] &\geq \sum_{j\in M}\sum_{i=1}^{n}\mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\max\left(\mathsf{SW}_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i}) - p'_j, 0\right)\cdot \mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\right] \\
&\quad \cdot \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\notin \texttt{SOLD}_i(\mathbf{v},\pi)\right] \\
&\geq \sum_{j\in M}\sum_{i=1}^{n}\mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\max\left(\mathsf{SW}_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i}) - p'_j, 0\right)\cdot \mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\right] \\
&\quad \cdot \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\notin \texttt{SOLD}(\mathbf{v},\pi)\right] \\
&\geq \sum_{j\in M}\mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\notin \texttt{SOLD}(\mathbf{v},\pi)\right]\sum_{i=1}^{n}\mathop{\mathbf{E}}_{\mathbf{v}}\left[\left(\mathsf{SW}_j(\mathbf{v}) - p'_j\right)\cdot \mathbb{1}\left[j\in X_i(\mathbf{v})\right]\right] \\
&= \sum_{j\in M}\mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\notin \texttt{SOLD}(\mathbf{v},\pi)\right]\cdot (p_j + (p_j - p'_j)). \qquad (5)
\end{aligned}$$

In the second inequality, we decreased each probability $\mathbf{Pr}[j\notin \texttt{SOLD}_i(\mathbf{v},\pi)]$ to $\mathbf{Pr}[j\notin \texttt{SOLD}(\mathbf{v},\pi)]$; the inequality holds as all the terms in the summation are non negative. In the third inequality we decreased the random variables under expectations and relabeled variables $(v_i,\widetilde{\mathbf{v}}_{\text{-}i})$ to $\mathbf{v}$. The last equality follows from (3). Inequality (5) is our desired bound on the sum of buyer utilities.

We now turn to the expected revenue, which is

$$\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathrm{Rev}(\mathbf{v},\pi)\right] = \sum_{j\in M}\mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right]\cdot (p_j - (p_j - p'_j)). \qquad (6)$$

Therefore, adding (5) and (6) we derive the following bound on the expected social welfare:

$$\begin{aligned}
\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} u_i(\mathbf{v})\right] + \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathrm{Rev}(\mathbf{v},\pi)\right] &\geq \left(\sum_{j\in M} p_j\right) + \sum_{j\in M}(p_j - p'_j)\cdot\left(1 - 2\mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right]\right) \\
&\geq \frac{1}{2}\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} v_i(X_i)\right] - \sum_{j\in M}|p_j - p'_j| \\
&\geq \frac{1}{2}\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} v_i(X_i)\right] - m\delta
\end{aligned}$$

as required. □

With Lemma 2.6 at hand, we are ready to complete the proof of Theorem 2.2.

*Proof.* (of Theorem 2.2) It remains to show how to compute an appropriate choice of prices $\mathbf{p}'$ satisfying the conditions of Lemma 2.6. Our approach will be to estimate $p_j = \frac{1}{2}\cdot \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}_j(\mathbf{v})]$ by repeatedly sampling a valuation profile $\hat{\mathbf{v}}\sim\mathcal{F}$ and computing $\frac{1}{2}\mathsf{SW}_j(\hat{\mathbf{v}})$. Since $\frac{1}{2}\mathsf{SW}_j(\hat{\mathbf{v}})$ is a random variable lying in $[0,1]$, standard concentration bounds imply that we can accurately estimate its expectation in a relatively small number $t$ of samples. To summarize, this procedure for computing $\mathbf{p}'$ is listed formally as Algorithm 2.

**ALGORITHM 2:** Price computation algorithm, paramaterized by positive integer $t$.

```
For each item j ∈ M:
  Repeat t times:
    Draw v ~ F, let X = A(v), and let i be the agent for which j ∈ X_i.
    Query the XOS oracle for v_i to find SW_j(v).
  Let p'_j be half of the average value of SW_j(v) seen over all t iterations.
return p'
```

We wish to choose $t$ large enough that, with probability at least $1-\varepsilon/n$, we will have $|p'_j - p_j| < \varepsilon/2m$ for all $j$. Fix any $j$ and note that $p'_j$ is the average of $t$ identical samples from a distribution supported on $[0,1]$, with expected value $p_j$. Thus, by the Hoeffding bound, we have that

$$\Pr[|p'_j - p_j| > \varepsilon/2m] < 2e^{-t(\varepsilon/2m)^2}.$$

We can therefore choose $t = (\log m + \log n - \log \varepsilon)4m^2/\varepsilon^2$ to get $\Pr[|p'_j - p_j| > \varepsilon/2m] < \varepsilon/mn$. Applying a union bound over all $j \in M$ we obtain a guarantee that $|p'_j - p_j| < \varepsilon/2m$ for all $j$ with probability at least $1-\varepsilon/n$, as desired.

Now, by setting $\delta = \varepsilon/2m$ in Lemma 2.6, we have that our computed prices generate welfare at least $\frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \varepsilon/2$, with probability at least $1-\varepsilon/n$. We conclude that our computed prices generate an expected welfare of at least

$$\left(\frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \frac{\varepsilon}{2}\right)\left(1-\frac{\varepsilon}{n}\right) > \frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \varepsilon,$$

as required. The last inequality follows since $\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] \le \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\sum_{i=1}^n v_i(M)] \le n$. ∎

## 2.3 A Lower Bound

A simple example shows that one cannot hope to find prices that achieve better than half of the optimal expected welfare, even for a single item. This is true even if the posted price mechanism can adaptively choose the order in which agents are considered, given the type realizations of agents who were considered previously (but not the type realizations of agents who have not yet been considered).

**Theorem 2.7.** *For any $\varepsilon > 0$, there exists a submodular combinatorial auction in which no posted price mechanism can guarantee expected welfare greater than $(\frac{1}{2}+\varepsilon)OPT$. This is true even for a single item, and even if the mechanism can adaptively choose the order in which buyers are considered.*

*Proof.* Suppose $\mathcal{F}$ is such that each agent has (large) value $X$ with probability $q = 1-(1-1/X)^{1/n}$ and value 1 otherwise. Note that $q = \theta(\frac{1}{Xn})$, and that the probability that *any* agent has value $X$ is exactly $1/X$. Thus the expected optimal welfare is $1\cdot(1-1/X) + X\cdot(1/X)$, which approaches 2 as $X$ grows large. On the other hand, no posted price $p$ obtains welfare greater than 1: if $p > X$ then no welfare is generated; if $p \in (1, X)$ then it generates welfare $X$ with probability $1/X$, for an expected welfare of 1; and if $p \le 1$, an arbitrary agent will buy (since the distributions are i.i.d., and hence all agents are indistinguishable prior to being considered). In the last case, the expected welfare is $1 + Xq = 1 + X\cdot\theta(\frac{1}{Xn}) = 1 + \theta(\frac{1}{n})$, which approaches 1 as $n$ grows large. Thus, for $X$ and $n$ sufficiently large, we conclude that no posted price can generate expected welfare greater than $(\frac{1}{2}+\varepsilon)$OPT. ∎

## 3 Posted Prices for General Valuations

A result similar to Theorem 2.2 also holds for $\mathcal{MPH}$-$k$ valuations, where we get $O(k)$-approximate DSIC mechanisms for functions with complementarity level $k$. We recall that $\mathcal{MPH}$-$k$ hierarchy covers valuation classes ranging from XOS functions (corresponding to $k = 1$) to general valuations (corresponding to $k = m$). Before stating the result formally, we will recall the definition of $\mathcal{MPH}$-$k$ from [19].

To formalize the *maximum over positive hypergraphs* ($\mathcal{MPH}$) hierarchy, we first need a few preliminaries. A hypergraph representation $w$ of valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is a set function that satisfies $v(S) = \sum_{T \subseteq S} w(T)$. Any valuation function $v$ admits a unique hypergraph representation and vice versa. A set $S$ such that $w(S) \neq 0$ is said to be a *hyperedge* of $w$. Pictorially, the hypergraph representation can be thought as a weighted hypergraph, where every vertex is associated with an item in $M$, and the weight of each hyperedge $e \subseteq M$ is $w(e)$. Then the value of the function for any set $S \subseteq M$, is the total value of all hyperedges that are contained in $S$.

The *rank* of a hypergraph representation $w$ is the cardinality $k$ of the largest hyperedge. The rank of $v$ is the rank of its corresponding $w$ and we refer to a valuation function $v$ with rank $k$ as a *hypergraph-k* valuation. If the hypergraph representation of $v$ is non-negative, i.e. for any $S \subseteq M$, $w(S) \geq 0$, then we refer to function $v$ as a *positive hypergraph-k* function (PH-$k$) [1].

We are now ready to present the class $\mathcal{MPH}$-$k$.

**Definition 3.1** (Maximum Over Positive Hypergraph-$k$ ($\mathcal{MPH}$-$k$) class [19])**.** *A monotone valuation function $v : 2^M \to \mathbb{R}^+$ is* Maximum over Positive Hypergraph-$k$ *($\mathcal{MPH}$-$k$) if it can be expressed as a maximum over a set of PH-k functions. That is, there exist PH-k functions $\{v_\ell\}_{\ell \in \mathcal{L}}$ such that for every set $S \subseteq M$,*

$$v(S) = \max_{\ell \in \mathcal{L}} v_\ell(S), \tag{7}$$

*where $\mathcal{L}$ is an arbitrary index set.*

The highest level of the hierarchy, $\mathcal{MPH}$-$m$ captures all monotone functions, and the lowest level, $\mathcal{MPH}$-1, captures all $\mathcal{XOS}$ functions.

Finally, we define what is meant by an $\mathcal{MPH}$-$k$ oracle, which is an extension of $\mathcal{XOS}$ oracles to higher levels of the $\mathcal{MPH}$ hierarchy. Suppose that valuation function $v$ is $\mathcal{MPH}$-$k$, with supporting PH-$k$ functions $\{v_\ell\}_{\ell \in \mathcal{L}}$. An $\mathcal{MPH}$-$k$-oracle for $v$ takes as input a set of items $S$, and returns the PH-$k$ function $v_\ell$ for which $v(S) = v_\ell(S)$. We will assume that this function $v_\ell$ is returned in its explicit hypergraph representation, i.e. as a list of weighted hyperedges. Note that the size of this representation depends on the number of hyperedges required to express the PH-$k$ functions $v_\ell$, and is at most $O(m^k)$. On a side note, it is this bound that leads to a runtime that is polynomial in $m^k$ in Theorem 3.1. If we restrict attention to MPH-$k$ valuations whose supporting PH-$k$ functions each have at most $r$ hyperedges, then this runtime dependency would change from $m^k$ to $r$.

Our result is cast in the following theorem, whose proof is deferred to Appendix A.

**Theorem 3.1.** *Suppose our Bayesian instance $\mathcal{F}$ over $\mathcal{MPH}$ valuations is given via a sample access to $\mathcal{F}$. Suppose that for every $\mathbf{v} \sim \mathcal{F}$ we have*

1. *black-box access to a welfare maximization algorithm $\mathcal{A}$ for combinatorial auctions,*
2. *an $\mathcal{MPH}$ query oracle for the valuations in the support of $\mathcal{F}$.*

*Then, for every $\varepsilon > 0$, one can compute item prices in time*[9] $\mathrm{POLY}(m^k, n, 1/\varepsilon)$ *that generate expected welfare of at least* $\frac{1}{4k}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[SW(\mathcal{A}(\mathbf{v}))] - \varepsilon$ *for any buyers' arrival order.*

The proof is similar to that of Theorem 2.2 for $\mathcal{XOS}$ buyers. However, there is an extra difficulty for $\mathcal{MPH}$-$k$ valuations: the notion of "contribution of an item to welfare" – used to determine an item's price in Theorem 2.2 – is not as straightforward as for $\mathcal{XOS}$ valuations. Thus one challenge in proving Theorem 3.1 is to appropriately account for the contributions of different items. The approach we take is to split the weight of each hyperedge evenly among its constituent items: the contribution of an item to the social welfare is the sum of that item's share of each hyperedge that contributes to the social welfare. With this definition in place, the proof closely follows that of Theorem 2.2. The details appear in Appendix A.

We also show that this result is essentially tight. Indeed, for each level $k$ of complementarities across the items in the $\mathcal{MPH}$-$k$ hierarchy, we may consider single minded (of size $k$) and unit-demand valuations. It turns out that item prices may result in an outcome with a linear (in the number of items) loss in social welfare. The example is deferred to Appendix B.

# 4 Discussion and Open Problems

We conclude with a few remarks. First, in our mechanisms, we consider an arbitrary order of arrivals, which may be chosen by an adversary after the prices are posted, but before the adversary observes the realization of the buyer valuations. It is not difficult to verify that the same results extend to an *adaptive* adversary, who chooses the arrival order sequentially; i.e., an adversary who observes which items have been purchased by previous buyers and even the realization of previous buyers' valuations, and chooses the next buyer to arrive based on this information. Our proof techniques (in Theorems 2.2 and 3.1) apply to this adaptive adversary as well.

Second, readers who are familiar with literature on Walrasian equilibrium will realize the similarities between the two models, but also the stark contrast. The main difference is whether agents arrive to the market *sequentially* (as in our model), or *simultaneously* (as in a Walrasian equilibrium). Recent results [21] have shown that in the simultaneous model (even when some items may remain unsold), there may be a *linear* loss in welfare for XOS buyers, even in a full information setting. Thus our work demonstrates a strong gap in welfare between simultaneous and sequential arrivals, when restricted to individual demand satisfaction.

Our model and results leave a number of directions for future research. First, the constant approximation for XOS valuations implies (by known results, see e.g. [5]) a logarithmic approximation for subadditive valuations. We conjecture that a constant approximation for subadditive valuations can be achieved.

Second, throughout the paper we assume that items are indivisible and heterogeneous. It would be interesting to partially relax these assumption. For example, one could assume that every item in the market has a few identical copies and that every buyer wants at most a single copy of each item. It would be interesting to analyze the efficiency of posted price mechanisms as a function of the minimal number of item copies. Given the negative results for valuations with high degree of complementarity, it would be particularly interesting to find relaxations that admit positive results, say for single minded buyers.

[9] The exponential dependence on $k$ in the runtime is related to the representation of $\mathcal{MPH}$ valuations. In particular, the output of an $\mathcal{MPH}$-$k$ oracle can be of size $O(m^k)$. One could reduce this bound by imposing constraints on the complexity of a valuation's $\mathcal{MPH}$ representation. This is discussed further in Appendix A.

## Acknowledgments

We thank David Kempe and Noam Nisan for helpful discussions regarding sequential sales. We thank Alon Eden and Amos Fiat for helpful discussions about arbitrary tie breaking. Michal Feldman is partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

## References

[1] Ittai Abraham, Moshe Babaioff, Shaddin Dughmi, and Tim Roughgarden. Combinatorial auctions with restricted complements. In *ACM Conference on Electronic Commerce*, pages 3–16, 2012.

[2] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*, pages 512–521, 2011.

[3] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In *ICS*, pages 166–178, 2011.

[4] X. Bei and Z. Huang. Bayesian incentive compatibility via fractional assignments. In *SODA*, 2011.

[5] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*, pages 700–709, 2011.

[6] Kshipra Bhawalkar and Tim Roughgarden. Simultaneous single-item auctions. In *WINE*, 2012.

[7] L. Blumrosen and N. Nisan. Combinatorial auctions. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 11, pages 267–299. Cambridge University Press, 2007.

[8] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *EC*, page 49, 2008.

[9] Liad Blumrosen and Noam Nisan. On the computational power of demand queries. *SIAM J. Comput.*, 39(4):1372–1391, 2009.

[10] Yang Cai and Christos H. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *ACM Conference on Economics and Computation*, pages 895–910, 2014.

[11] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce*, pages 243–251, 2007.

[12] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010.

[13] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The Power of Randomness in Bayesian Optimal Mechanism Design. In *the 11th ACM Conference on Electronic Commerce (EC)*, 2010.

[14] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *ICALP*, pages 820–832, 2008.

[15] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *APPROX-RANDOM*, pages 89–103, 2007.

[16] Shahar Dobzinski and Jan Vondrák. The computational complexity of truthfulness in combinatorial auctions. In *EC*, pages 405–422, 2012.

[17] Shaddin Dughmi and Jan Vondrák. Limitations of randomized mechanisms for combinatorial auctions. In *FOCS*, pages 502–511, 2011.

[18] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142, 2009.

[19] Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. A unified approach to valuations with restricted complements. In *Working paper*, 2014.

[20] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *STOC*, pages 201–210, 2013.

[21] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial walrasian equilibrium. In *STOC*, pages 61–70, 2013.

[22] Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. Bayesian incentive compatibility via matchings. In *SODA*, pages 734–747, 2011.

[23] Jason D. Hartline and Brendan Lucier. Bayesian algorithmic mechanism design. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 301–310, 2010.

[24] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *EC*, pages 295–296, 2011.

[25] Theodore P. Hill and Robert P. Kertz. A survey of prophet inequalities in optimal stopping theory. *CONTEMPORARY MATHEMATICS*, 125, 1992.

[26] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th Symposium on Theory of Computing Conference*, pages 123–136, 2012.

[27] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270 – 296, 2006.

[28] Shengwu Li. Obviously strategy-proof mechanisms. In *Working paper*.

[29] Renato Paes Leme. Gross substitutability: An algorithmic survey. 2013.

[30] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC*, pages 67–74, 2008.

# APPENDIX

## A $\mathcal{MPH}$ valuations

This section is dedicated to the proof of Theorem 3.1.

**Proof of Theorem 3.1:**

We closely follow the proof of Theorem 2.2 for $\mathcal{XOS}$ buyers. However, there is an extra difficulty for $\mathcal{MPH}$-$k$ valuations as "contribution of an item to welfare" is not as straightforward as for $\mathcal{XOS}$ valuations. Our main new challenge will be to appropriately account for the contributions of different items. We first describe an ideal price vector $\mathbf{p}$ which we would like to use for the distribution $\mathcal{F}$. For each fixed valuation profile $\mathbf{v} = (v_1, \ldots, v_n)$ we consider allocation $\mathbf{X}(\mathbf{v}) = (X_1(\mathbf{v}), \ldots, X_n(\mathbf{v}))$ returned by black-box algorithm $\mathcal{A}$. For each MPH-$k$ valuation function $v_i$ we take the respective PH-$k$ function $a_i$ for the set $X_i(\mathbf{v})$, i.e., $v_i(S) \geq a_i(S)$ for any $S \subset [m]$ and $v_i(X_i(\mathbf{v})) = a_i(X_i(\mathbf{v}))$. Write $w_i$ for the hypergraph representation of $a_i$; then, by definition, $a_i(S) = \sum_{T \subseteq S} w_i(T)$ for all $S \subseteq X_i(\mathbf{v})$.

For every $\mathbf{v}$, every buyer $i$, and each item $j \in X_i(\mathbf{v})$, we define

$$p_j(\mathbf{v}) = \frac{1}{\alpha} \sum_{\substack{T \ni j \\ T \subseteq X_i(\mathbf{v})}} \frac{w_i(T)}{|T|},$$

where $\alpha$ is a constant to be determined later. The price $p_j(\mathbf{v})$ for item $j$ is the ideal price we would like to set in the full-information setting, if we knew the valuation profile $\mathbf{v}$.

We can now define an ideal price of item $j$ in the Bayesian setting, which will be

$$p_j = \mathop{\mathbf{E}}_{\mathbf{v} \sim \mathcal{F}} \left[ p_j(\mathbf{v}) \right].$$

The following Lemma relates the full-information prices for a subset of items to the marginal impact on a buyer's value if those items are removed from an allocation.

**Lemma A.1.** *For any $\mathbf{v}$, any buyer $i$, and any $Q \subseteq X_i(\mathbf{v})$,*

$$v_i(X_i(\mathbf{v}) \backslash Q) + \alpha k \cdot \sum_{j \in Q} p_j(\mathbf{v}) \geq v_i(X_i(\mathbf{v})).$$

*Proof.*

$$\begin{aligned}
v_i(X_i(\mathbf{v}) \backslash Q) + \alpha k \cdot \sum_{j \in Q} p_j(\mathbf{v}) &= \sum_{T \subseteq X_i(\mathbf{v}) \backslash Q} w_i(T) + \alpha k \cdot \sum_{j \in Q} \frac{1}{\alpha} \cdot \sum_{\substack{T \ni j \\ T \subseteq X_i(\mathbf{v})}} \frac{w_i(T)}{|T|} \\
&\geq \sum_{T \subseteq X_i(\mathbf{v}) \backslash Q} w_i(T) + \sum_{j \in Q} \sum_{\substack{T \ni j \\ T \subseteq X_i(\mathbf{v})}} w_i(T) \\
&\geq \sum_{T \subseteq X_i(\mathbf{v}) \backslash Q} w_i(T) + \sum_{\substack{T \subseteq X_i(\mathbf{v}) \\ T \cap Q \neq \emptyset}} w_i(T) \\
&= \sum_{T \subseteq X_i(\mathbf{v})} w_i(T) \\
&= v_i(X_i(\mathbf{v}))
\end{aligned}$$

where the first inequality follows because $w_i(T) > 0$ only for $T$ with $|T| \leq k$, and the second inequality follows by noting that each hyperedge $T$ counted in the second summation must have a non-empty intersection with $Q$ and is counted $|T \cap Q| \geq 1$ times. $\square$

The next Lemma estimates the expected social welfare of a mechanism with posted prices that are close to the ideal $\mathbf{p}$.

**Lemma A.2.** *Let $\mathbf{p}'$ be any price vector such that $|p'_j - p_j| < \delta$ for all $j$. Then, any greedy consumption under prices $\mathbf{p}'$ results in expected welfare of at least $\frac{1}{4k}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathit{SW}(\mathcal{A}(\mathbf{v}))] - 2m\delta$.*

*Proof.* Given the prices $\mathbf{p}'$, let $\pi$ be the (adversarial) order of arrival. We are going to bound the sum of buyers' utilities in expectation over $\mathcal{F}$. To do so for each fixed $i$ and $\mathbf{v} = (v_i, \mathbf{v}_{\text{-}i})$ we consider another random valuation profile $\widetilde{\mathbf{v}}_{\text{-}i} \sim \mathcal{F}_{\text{-}i}$ independent of $\mathbf{v}$ and allocation $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ returned by $\mathcal{A}$ on the valuation profile $(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$. Let $S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i}) := X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) \cap \mathtt{SOLD}_i(\mathbf{v}, \pi)$ be the subset of items in $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ that are already sold when buyer $i$ is selected to make a purchase. Let $R_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i}) := X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) \setminus \mathtt{SOLD}_i(\mathbf{v}, \pi)$ be the subset of items in $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ that remain at this time. We note that buyer $i$ could have picked the set $R_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})$ and, therefore, his utility is at least the utility he would get from this set. Thus we have

$$u_i(\mathbf{v}) \geq \mathop{\mathbf{E}}_{\widetilde{\mathbf{v}}_{\text{-}i}}\left[v_i(R_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})) - \sum_{j \in R_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})} p'_j\right]$$
$$\geq \mathop{\mathbf{E}}_{\widetilde{\mathbf{v}}_{\text{-}i}}\left[v_i(R_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})) - \sum_{j \in X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})} p'_j\right].$$

Applying Lemma A.1 to valuation profile $(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ and set $Q = \mathtt{SOLD}_i(\mathbf{v}, \pi)$, we conclude

$$u_i(\mathbf{v}) \geq \mathop{\mathbf{E}}_{\widetilde{\mathbf{v}}_{\text{-}i}}\left[v_i(X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})) - \alpha k \sum_{j \in S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})} p_j(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) - \sum_{j \in X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})} p'_j\right].$$

We now sum over all $i$ and take an expectation over $\mathbf{v} \sim \mathcal{F}$ to conclude that

$$\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right] \geq \sum_i \mathop{\mathbf{E}}_{\mathbf{v}, \widetilde{\mathbf{v}}_{\text{-}i}}\left[v_i(X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})) - \sum_{j \in X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})} p'_j\right]$$
$$- \alpha k \cdot \sum_i \mathop{\mathbf{E}}_{\mathbf{v}, \widetilde{\mathbf{v}}_{\text{-}i}}\left[\sum_{j \in S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})} p_j(v_i, \widetilde{\mathbf{v}}_{\text{-}i})\right]. \tag{8}$$

Let us analyze separately the two summations on the RHS of (8). For the first summation, note that $\mathbf{v}_{\text{-}i}$ does not appear in the expression within the expectation. Thus, by applying a change of variables, we have

$$\begin{aligned}
\sum_i \mathop{\mathbf{E}}_{\mathbf{v},\widetilde{\mathbf{v}}_{\text{-}i}}\left[v_i(X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})) - \sum_{j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})} p'_j\right] &= \sum_i \mathop{\mathbf{E}}_{\mathbf{v}}\left[v_i(X_i(\mathbf{v})) - \sum_{j\in X_i(\mathbf{v})} p'_j\right] \\
&= \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \sum_j p'_j \\
&\geq \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \sum_j p_j - \delta m \\
&= \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \frac{1}{\alpha}\cdot\sum_j \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i \sum_{\substack{T\ni j\\ T\subseteq X_i(\mathbf{v})}} \frac{w_i(T)}{|T|}\right] - \delta m \\
&= \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \frac{1}{\alpha}\cdot\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \delta m \\
&= \left(1-\frac{1}{\alpha}\right)\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \delta m. \qquad (9)
\end{aligned}$$

For the second summation in RHS of (8), we first recall that sets $\texttt{SOLD}_i$ and $S_i(v_i,\mathbf{v}_{\text{-}i},\widetilde{\mathbf{v}}_{\text{-}i})$ are defined for the prices $\mathbf{p}'$. Further note that since $\mathbf{v}_{\text{-}i}$ and $\widetilde{\mathbf{v}}_{\text{-}i}$ are drawn independently, we have

$$\begin{aligned}
&\sum_i \mathop{\mathbf{E}}_{\mathbf{v},\widetilde{\mathbf{v}}_{\text{-}i}}\left[\sum_{j\in S_i(v_i,\mathbf{v}_{\text{-}i},\widetilde{\mathbf{v}}_{\text{-}i})} p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right] \\
&= \sum_i \mathop{\mathbf{E}}_{\mathbf{v},\widetilde{\mathbf{v}}_{\text{-}i}}\left[\sum_j \mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\cdot \mathbb{1}\left[j\in \texttt{SOLD}_i(\mathbf{v}_{\text{-}i},\pi)\right]\cdot p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right] \\
&= \sum_j\sum_i \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\cdot p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\mathop{\mathbf{Pr}}_{\mathbf{v}_{\text{-}i}}\left[j\in \texttt{SOLD}_i(\mathbf{v}_{\text{-}i},\pi)\right] \\
&\leq \sum_j\sum_i \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\cdot p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right] \\
&= \sum_j \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right]\sum_i \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\cdot p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right] \\
&= \sum_j \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right]\cdot p_j \leq \sum_j \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right]\cdot p'_j + \delta m \\
&= \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathrm{Rev}(\mathbf{v},\pi)\right] + \delta m. \qquad (10)
\end{aligned}$$

Substituting (9) and (10) into (8), we have

$$\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right] \geq \left(1-\frac{1}{\alpha}\right)\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \delta m - \alpha k\cdot \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathrm{Rev}(\mathbf{v},\pi)\right] - \alpha k\delta m. \qquad (11)$$

As long as $\alpha k \geq 1$, we can rearrange and conclude

$$\alpha k \left( \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}} \left[ \sum_{i=1}^{n} u_i(\mathbf{v}) \right] + \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}} \left[ \text{Rev}(\mathbf{v}, \pi) \right] \right) \geq \left(1 - \frac{1}{\alpha}\right) \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}} \left[ \sum_{i=1}^{n} v_i(X_i) \right] - 2\alpha k \delta m.$$

Taking $\alpha = 2$, we conclude that the expected welfare of the Posted Pricing Mechanism is within a factor $4k$ of the expected welfare of $\mathcal{A}$ and small additive error of $2m\delta$, as required. □

We continue with the proof of Theorem 3.1. Following the same analysis as in Theorem 2.2 for each item $j$ we can estimate the price $p'_j$ by sampling $t = (\log m + \log n - \log \varepsilon)16m^2/\varepsilon^2$ valuation profiles, so that $\Pr[|p'_j - p_j| > \varepsilon/4m] < \varepsilon/mn$. We compute $p_j{}'$ for each sample (using algorithm $\mathcal{A}$ and the MPH-$k$ query oracle) and take the average of all prices seen. Applying a union bound over all $j \in M$ we obtain a guarantee that $|p'_j - p_j| < \varepsilon/4m$ for all $j$ with probability at least $1 - \varepsilon/n$. Now, by setting $\delta = \varepsilon/4m$ in Lemma A.2 we have our computed prices $\mathbf{p'}$ to generate welfare of at least $\frac{1}{4k}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \frac{\varepsilon}{2}$ with probability at least $1 - \varepsilon/n$.

Finally, we conclude that generated expected welfare is at least

$$\left(\frac{1}{4k}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \frac{\varepsilon}{2}\right)\left(1 - \frac{\varepsilon}{n}\right) > \frac{1}{4k}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \varepsilon,$$

as required. The last inequality follows, since $\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] \leq \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\sum_{i=1}^{n} v_i(M)] \leq n$.

# B Lower bound for $\mathcal{MPH}$

**Theorem B.1.** *For any $\varepsilon > 0$, there exists a combinatorial auction in which no posted price mechanism can guarantee expected welfare greater than $(\frac{1}{m-\varepsilon})OPT$.*

*Proof.* Suppose there are $m$ identical items in the market and two buyers. Let the first buyer have unit-demand valuation 1 per item and the second single-minded buyer have value $m - \varepsilon$ for the set of all $m$ items and 0 value for any smaller subset. The optimal social welfare OPT is $m - \varepsilon$, where the second buyer is allocated all $m$ items.

Let the seller fix prices on the items, and let $p$ be the price of the cheapest item. We let the first buyer arrive first. If $p < 1$, he will buy one of the cheapest items. Then the second buyer has 0 value for the remaining items, which results in a social welfare of 1. If $p \geq 1$, then regardless of whether the first buyer purchases the second buyer will purchase nothing, as he derives value $m - \varepsilon$ from the entire set, which costs at least $m$. Therefore, the social welfare in the latter case is also at most 1. We conclude that the social welfare does not exceed 1 in either of the cases, which gives us the claimed linear gap of $m - \varepsilon$ with respect to the optimal social welfare. □