Combinatorial Auctions via Posted Prices

Michal Feldman* Nick Gravin† Brendan Lucier‡

Abstract

We study anonymous posted price mechanisms for combinatorial auctions in a Bayesian framework. In a posted price mechanism, item prices are posted, then the consumers approach the seller sequentially in an arbitrary order, each purchasing her favorite bundle from among the unsold items at the posted prices. These mechanisms are simple, transparent and trivially dominant strategy incentive compatible (DSIC).

We show that when agent preferences are fractionally subadditive (which includes all submodular functions), there always exist prices that, in expectation, obtain at least half of the optimal welfare. Our result is constructive: given black-box access to a combinatorial auction algorithm A, sample access to the prior distribution, and appropriate query access to the sampled valuations, one can compute, in polytime, prices that guarantee at least half of the expected welfare of A. As a corollary, we obtain the first polytime (in n and m) constant-factor DSIC mechanism for Bayesian submodular combinatorial auctions, given access to demand query oracles. Our results also extend to valuations with complements, where the approximation factor degrades linearly with the level of complementarity.

1 Introduction

The canonical problem in market design is to efficiently allocate a set of m resources among a set of n 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 AGT is to marry algorithmic solutions to such problems with the economic principles that underpin market design.

As an example, take the problem of designing an incentive compatible mechanism for combinatorial auctions with submodular bidders. The underlying optimization problem is NP-hard, and simple greedy methods achieve a constant approximation. On the other hand, whether there is a truthful polytime constant-factor approximation has remained a vexing and major open question for over a decade[1]. One might therefore wonder whether it is even feasible to implement a combinatorial auction mechanism that is both computationally tractable and economically appealing.

When designing an economically viable mechanism, it is desirable to not only respect incentives and computational constraints, but also to resolve allocation decisions in a straightforward and transparent fashion. For example, a simple and natural approach is to resolve a market using *posted prices*. One might imagine an implementation in which items are assigned *anonymous* prices (i.e., all agents face the exact same prices), then agents arrive to the market and each consumes his most-desired bundle under the given prices. Such a methodology exhibits many desirable properties: it is simple, decentralized, and easy to implement. It is also transparent in the sense that a buyer does not need to understand the price-setting method in order to understand how to participate — he simply behaves as a price-taker and consumes his preferred bundle. Therefore, it is also trivially incentive compatible.

Posted price mechanisms are highly applicable when markets are large and the aggregate demands of buyers can be accurately predicted. For instance, if the buyers' valuations are public knowledge and satisfy the gross substitutes condition, then there always exist prices that efficiently clear the market [23, 29]. Similar results hold in large markets for arbitrary valuations [32]. However, if valuations are not fully known to the seller and are unpredictable, as in the submodular combinatorial auction problem described above, then it is unclear

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

[1]For some models of valuation access, such as the value query model, it is in fact known that no sub-polynomial worst-case approximation is possible [17].

how reasonable prices can be set. This motivates the approach of implementing an optimization algorithm as an *auction*, where buyers submit competing bids that are treated as input. As Milgrom writes, "When goods are not standardized or when the market clearing prices are highly unstable, posted prices work poorly, and auctions are usually preferred" [33].

We are left with a dichotomy. Posted prices form a simple and natural market instrument, but they generate efficient outcomes only under special circumstances and full information over the buyers' preferences. General auctions are more widely applicable, but can be significantly more complex to execute and participate in; moreover, we still do not have reasonable auction designs for many allocation problems of interest. A natural question arises:

> to what extent can *approximately* efficient outcomes be implemented using anonymous posted prices in settings of incomplete information?

In this work we study the power of anonymous posted prices in Bayesian settings, where the designer knows the distribution over the agents' valuations but not their realizations. The Bayesian setting imposes additional challenges over the full information setting typical in the market equilibrium literature. In particular, readers familiar with this literature will note that, in this probabilistic setting, there do not exist prices that are guaranteed to satisfy all agent demands simultaneously. To mitigate this problem, we consider posted price mechanisms that admit the agents sequentially, in an arbitrary order, to select their most preferred bundle from the remaining items.

We devise polytime posted price mechanisms, for several classes of valuations, that achieve nearly optimal social welfare given appropriate access to the distribution of the agents' valuations. Specifically, for XOS valuations (a strict superset of submodular valuations) we devise a mechanism that obtains a constant fraction of the optimal social welfare. Notably, this implies that we can obtain an $O(1)$-approximate dominant strategy incentive compatible mechanism for Bayesian XOS valuations, whose running time is polynomial in n and m given access to demand queries. Prior to this work, the best-known polytime mechanisms (posted-price or otherwise) were either polylogarithmic approximations [15] or had runtimes that were polynomial in the support size of an agent's valuation distribution, which could be exponential in n and m [4, 24]. In addition, for general monotone valuations, we devise a mechanism that obtains an approximation factor that degrades gracefully with the level of complementarity of the functions, as captured by the maximum-over-hypergraph (MPH) hierarchy, recently introduced in [20]. A more detailed description of our results, along with a comparison to the related literature, appears below.

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 $v_i(\cdot) : 2^M \to \mathbb{R}_{\geq 0}$ 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 $[0, 1]$. 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, for every buyer $i \in [n]$, and $X_i \cap X_k = \emptyset$ 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 $\mathsf{SW}(\mathbf{X}) = \sum_{i=1}^{n} v_i(X_i)$, and the optimal welfare is denoted by OPT.

The utility of buyer i being allocated bundle X_i under prices $\mathbf{p}$ is $u_i(X_i, \mathbf{p}) = v_i(X_i) - \sum_{j \in X_i} p_j$. Given prices $\mathbf{p} = (p_1, \ldots, p_m)$, the *demand correspondence* $D_i(M, \mathbf{p})$ of buyer i contains the sets of objects that maximize buyer i's utility.

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 α-approximation (for social welfare) if

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

1.2 Anonymous Posted Price Mechanisms The posted price mechanisms considered in this paper proceed in the following steps:

1. (pricing phase) a price vector $\mathbf{p}$ is determined, based on $\mathcal{F}$;
2. (arrival order phase) an arrival order π is determined;
3. (value realization phase) A value $v_i \sim \mathcal{F}_i$ is realized for every buyer i, known only to buyer i;
4. (consumption phase) The buyers arrive to the market according to the order π, and each buyer receives his most desired bundle among all remaining items. That is, for every buyer i, $X_i \in D_i(\mathbf{p}, M \setminus \bigcup_{j <_\pi i} X_j)$. Buyer i pays the sum of item prices in his bundle X_i.

We refer to these as anonymous posted price mechanisms, though sometimes we will leave out "anonymous" for brevity. We note that the mechanism can be described in two equivalent ways, namely as an indirect or

a direct mechanism. In both versions, the mechanism sets the prices and determines[2] the arrival order. The difference is that in the indirect implementation, the agents arrive and purchase their desired items in the determined order, whereas in the direct revelation implementation, the buyers report their valuations, and the mechanism simulates the consumption phase on their behalf. This latter simulation requires that the mechanism have access to demand oracles[3].

It is easy to verify that any posted price mechanism that adheres to this structure is trivially dominant strategy incentive compatible (DSIC). Moreover, it is clearly *weakly group strategyproof*, meaning that no coalition of agents can deviate in a way that strictly benefits each one of them.

All of our results hold regardless of the order selected in the arrival order phase, so we can think of the arrival order as being chosen adversarially. In the direct revelation implementation it would be natural to imbue the mechanism with the power to choose the ordering, but we view obliviousness to the order as increasing the robustness of the mechanism.

Beyond the strong incentive characteristics of posted price mechanisms, they are appealing in their simplicity. In particular, they proceed by setting a single price vector and using it for all agents.

1.3 Our results We establish welfare guarantees for XOS (i.e., fractionally subadditive) valuations (a strict superset of submodular valuations), and for MPH-k valuations — a hierarchy that spans all monotone valuations, parameterized by the complementarity level k. Our main contribution is the following:

Theorem: [2-approximation for XOS] Given black-box access to an algorithm $\mathcal{A}$ for XOS valuations, sample access to $\mathcal{XOS}$ distributions $\mathcal{F}$, an $\mathcal{XOS}$ query oracle for the valuations in the support of $\mathcal{F}$, and a demand oracle for the valuations, there exists a posted price mechanism that, for every XOS valuation profile $\mathbf{v}$, and every ε, returns an outcome that gives expected social welfare of at least $\frac{1}{2}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathcal{A}(\mathbf{v})] - \varepsilon$, and runs in time $\text{POLY}(n, m, 1/\varepsilon)$.

[2] All of our results hold under arbitrary arrival orders, so one can also think of an adversary as choosing the order.

[3] While answering a demand query might be NP-hard in some cases, in the context of combinatorial auctions it is natural to expect the agent to be able to answer a demand query. Otherwise it would be unreasonable to expect any mechanism to satisfy the agents' demands. Furthermore, in some cases there are polytime algorithms for answering demand queries. For example, for gross substitutes valuations demand queries can be implemented with a polynomial number of value queries [35].

We note that the factor of 2 is tight (even for a single item), as established by the illustrating example in the end of this section. We also note that the demand oracle requirement is needed only for simulating agent behavior in the consumption phase. This theorem implies a set of results for XOS valuations and special cases thereof, including submodular and gross substitutes valuations. Before presenting the corollaries, we wish to emphasize the strengths of this result in light of the previous literature:

- **(Truly) polytime:** Many previous mechanisms have (essentially) pseudo-polynomial runtime, in the sense that they are polynomial in the size of the agents' type space, which may plausibly be exponential in m and n. The runtime of our mechanisms is polynomial in m and n, independent of the type space sizes.
- **DSIC:** Many existing mechanisms for settings of incomplete information exhibit the weaker notion of Bayesian incentive compatibility (BIC). Our mechanism exhibits the stronger notion of DSIC, and is also weakly group strategyproof.

The heart of the construction of our mechanism lies in the appropriate price assignment; this is the part where oracle access is needed (except for the consumption phase that requires access to demand oracles). Together with known algorithmic results, our theorem implies the following polytime approximation results for XOS, submodular and gross substitutes valuations.

Theorem: [Computational results] We devise polytime (in n and m) DSIC mechanisms with the following guarantees:

- **XOS:** A $\frac{e}{2(e-1)}$-approximation, given sample access to the distribution, value and demand oracles, and XOS query oracles.
- **Submodular:** A $\frac{e}{2(e-1)}$-approximation, given sample access to the distribution, value oracles, and demand oracles. Demand oracles are required only in the consumption phase.
- **Gross substitutes (GS):** A $\frac{1}{2}$-approximation, given only sample access to the distribution and value oracles.

So far we have only considered complement-free valuations. Our second main result concerns valuation functions that exhibit complementarities.

Theorem: [k-approximation for MPH−k] Given a black-box access to an algorithm $\mathcal{A}$ for MPH−k valuations, a sample access to MPH−k distributions $\mathcal{F}$, an MPH−k query oracle for the valuations in the

support of $\mathcal{F}$, and a demand oracle for the valuations, there exists a posted price mechanism that, for every MPH−k valuation profile $\mathbf{v}$, and every ε, returns an outcome that gives expected social welfare of at least $\frac{1}{k}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathcal{A}(\mathbf{v})] - \varepsilon$, and 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, the last theorem implies the existence of posted prices that gives expected social welfare of at least $\frac{1}{\Omega(m)}$ fraction of OPT. We also show that this bound is tight.

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[4] with value greater than p, if any exists?

It is known that one cannot hope to achieve better than half of the optimal expected welfare [28]. For instance, suppose F is such that each agent has (large) value X with probability $q = 1 - (1 - 1/X)^{1/n}$ and value 1 otherwise. Then the probability that *any* agent has value X is $1/X$, and hence 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 obtains welfare greater than 1: if $p > 1$, it generates welfare X with probability $1/X$; whereas if $p \leq 1$, an arbitrary agent will buy and the expected welfare is $1 + Xq = 1 + O(\frac{1}{n})$. A posted price therefore cannot extract more than half of OPT.

However, there is a simple pricing scheme for a single item that yields half of the optimal social welfare; such methods are known from the prophet inequality literature [30]. Specifically, set price p equal to half of the expected highest value. To see why this works, write π 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 revenue from the auction is precisely $p \cdot \pi = \frac{1}{2}\mathbb{E}[v_{i^*}] \cdot \pi$. On the other hand, the probability that nobody buys the item ahead of buyer i^* is at least $1 - \pi$. The expected surplus (value minus payments) of buyer i^*, conditioned on item being still available, is at least $\mathbb{E}[v_{i^*}] - p = \frac{1}{2}\mathbb{E}[v_{i^*}]$. Putting this together, we have that the expected welfare of the auction, which equals the expected revenue plus the expected buyer surplus, is at least $\frac{1}{2}\mathbb{E}[v_{i^*}] \cdot \pi + \frac{1}{2}\mathbb{E}[v_{i^*}] \cdot (1 - \pi) = \frac{1}{2}\mathbb{E}[v_{i^*}]$, as claimed.

Our main result shows that the reasoning given in the single-item example can be extended to markets with multiple heterogeneous items for sale and asymmetric buyers, as long as buyer preferences lie in the class of XOS valuations.

1.4 Related work Our work is part of the recent body of literature on simple, non-optimal mechanisms [26, 14]. The design and performance of simple mechanisms is an active field of research, motivated by the observation that in practice, designers are often willing to trade truthfulness or optimality for simplicity. Canonical examples include the generalized second price (GSP) auctions for online advertising [18, 36], and the ascending price auction for electromagnetic spectrum allocation [34].

A particularly relevant example of a simple mechanism is the *simultaneous item* auction, in which agents make simultaneous, separate bids on multiple items. Such auctions are not truthful, but achieve, at equilibrium, a constant approximation to the optimal welfare when agents have complement-free valuations [6, 13, 21, 27]. A conclusion from this is that by restricting bidders to bid only on individual items, rather than on packages as in the VCG auction, one loses only a constant factor from the optimal welfare. Our results show that it is possible to go one step further, in terms of simplicity. Indeed, posted price mechanisms not only handle items separately but also forego competition entirely and simply publish prices on individual items, and simultaneously exhibit strong incentive compatibility.

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 form of posted-price mechanism extracts a constant fraction of the optimal revenue [10, 11, 12]. When there is only a single item for sale, this constant factor persists even when the distributions are unknown, as long as they are MHR [3]. Notably, these works all apply a more relaxed notion of a posted-price mechanism, in which one can set different prices for each customer. In contrast, the mechanisms we consider use non-discriminatory pricing.

Our work relates to the design of truthful submodular combinatorial auctions. 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. It is known that no sub-polynomial factor is possible under the value query

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

model [17], or for succinctly-described valuations [16]. 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 [25], Hartline et al. [24] 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. In particular, one can apply the latter two to constant-factor approximations for submodular CAs to obtain BIC constant-factor mechanisms. However, these mechanisms 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[5].

2 Preliminaries

Valuation Classes We study both complement-free valuations, and valuations that exhibit complementarities. There is a standard hierarchy of complement-free valuations (see [31]): additive $\subset$ gross substitutes $\subset$ submodular $\subset$ XOS $\subset$ subadditive.

additive $v(S) = \sum_{j \in S} v(\{j\})$ for all $S \subset M$.

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 $S \subseteq M$, $v(S) = \max_{1 \leq i \leq k} A_i(S)$.

subadditive for any subsets $S_1, S_2 \subset M$, $v(S_1) + v(S_2) \geq v(S_1 \cup S_2)$.

To study valuations with complements, we consider the hierarchy *maximum over positive hypergraphs* **(MPH)**, introduced recently by [20]. 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 4.

[5] 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.

Computational model 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 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) $v_i(T) = A_i(T)$;

While value oracle is the least computationally demanding for the buyers and the seller, the demand oracle captures the most basic decision problem a buyer faces in a market with item prices. Since the primary focus of this paper is on the pricing mechanisms, we assume throughout the paper an access to demand and value oracles for granted. We note that XOS oracles are less commonly used in the literature. However, for some classes of valuations XOS oracle can be implemented via polynomially many queries to value oracle, e.g., for any submodular function.

3 Posted Prices for XOS Valuations

The main theorem in this section is the following.

THEOREM 3.1. *Let distribution $\mathcal{F}$ over $\mathcal{XOS}$ valuation profiles be 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 $\mathcal{XOS}$ query oracle (for valuations sampled from $\mathcal{F}$).*

Then, for any $\varepsilon > 0$, we can compute item prices in $\mathrm{POLY}(m, n, 1/\varepsilon)$ *time 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 3.1, let us discuss some implications. First, note that using an α-approximation algorithm for $\mathcal{A}$ in Theorem 3.1 results in a posted price mechanism with approximation factor $\alpha/2$, minus an additive error term that can be made

as small as desired. Recall that implementing the consumption phase in a direct revelation mechanism does require access to demand queries; note that Theorem 3.1 and its corollaries below refer specifically to the pricing phase.

If we assume access to demand oracles, then we can use the polytime algorithm of Feige [19] with approximation factor $1-1/e$ as a black box. Theorem 3.1 then implies the existence of a DSIC mechanism with expected social welfare at least $\frac{e}{2(e-1)}\text{OPT}-\varepsilon$ and runtime $\text{POLY}(m,n,1/\varepsilon)$.

For submodular valuations, one could instead use the algorithm by Vondrak [37] 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 3.1. *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 $\frac{e}{2(e-1)}\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))]-\varepsilon$.*

As before, one can implement the mechanism from Corollary 3.1 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 [35], and an optimal allocation can be computed in polynomial time using demand queries [7]. The following corollary follows:

COROLLARY 3.2. *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 3.1 can be implemented as a direct revelation mechanism, using only value queries.

Proof of Theorem 3.1. We now proceed with the proof of Theorem 3.1. The proof will proceed in two parts. We begin with Lemma 3.1, which establishes the *existence* of prices that achieve the desired welfare properties, without regard for computation. In fact, Lemma 3.1 will also establish 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 3.1.

Before delving into the details of the proof, we need the following definition of an item's welfare contribution. Fix a valuation profile $\mathbf{v}=(v_1,\dots,v_n)$ and algorithm $\mathcal{A}$, and let $\mathbf{X}=(X_1,\dots,X_n)$ be the allocation $\mathcal{A}(\mathbf{v})$. For each XOS valuation function $v_i(\cdot)$, define the *corresponding additive representative* function for the set X_i as the function $A_i(\cdot)$ 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 3.1. *Given a distribution $\mathcal{F}$ over $\mathcal{XOS}$ valuations, let $\mathbf{p}$ be the price vector defined as*

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

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

Proof. First, by the definition of p_j,

$$p'_j = \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathsf{SW}_j(\mathbf{v})-p'_j\right]+2(p'_j-p_j) \tag{3.1}$$

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

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 $\texttt{SOLD}_i(\mathbf{v},\pi)$ denote the set of items that have been sold before the arrival of buyer i. Recall that buyer i picks an allocation[6] that maximizes his utility with respect to his valuation v_i and prices $\mathbf{p}$, from among the items in $M\setminus\texttt{SOLD}_i(\mathbf{v},\pi)$.

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 $S_i(v_i,\mathbf{v}_{\text{-}i},\widetilde{\mathbf{v}}_{\text{-}i}) := X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\setminus\texttt{SOLD}_i(\mathbf{v},\pi)$ 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

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

[6] Note that if a buyer has more than one bundle in his demand correspondence, then we assume that ties can be broken arbitrarily – even adversarially.

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

$$\begin{aligned}\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} u_i(\mathbf{v})\right] \geq \sum_{j\in M}\sum_{i=1}^{n} \mathop{\mathbf{E}}_{\substack{v_i,\widetilde{\mathbf{v}}_{-i},\\ \widetilde{\mathbf{v}}_{-i}}}\Big[\mathbb{1}\left[j \in X_i(v_i,\widetilde{\mathbf{v}}_{-i})\right] \\ \cdot \max\left(\mathsf{SW}_j(v_i,\widetilde{\mathbf{v}}_{-i}) - p'_j, 0\right) \\ \cdot \mathbb{1}\left[j \notin \mathtt{SOLD}_i(\mathbf{v},\pi)\right]\Big]. \qquad (3.2)\end{aligned}$$

We further observe that $\mathtt{SOLD}_i(\mathbf{v},\pi)$ does not depend on v_i. That is, $\mathtt{SOLD}_i(\mathbf{v},\pi) = \mathtt{SOLD}_i(\mathbf{v}_{-i},\pi)$. Therefore, we can rewrite (3.2) as follows:

$$\begin{aligned}\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} u_i(\mathbf{v})\right] \geq \sum_{j\in M}\sum_{i=1}^{n} \mathop{\mathbf{P}}_{\mathbf{v}}\mathrm{r}\left[j \notin \mathtt{SOLD}_i(\mathbf{v},\pi)\right] \\ \cdot \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{-i}}\left[\max\left(\mathsf{SW}_j(v_i,\widetilde{\mathbf{v}}_{-i}) - p'_j, 0\right)\cdot \mathbb{1}\left[j \in X_i(v_i,\widetilde{\mathbf{v}}_{-i})\right]\right] \\ \geq \sum_{j\in M}\sum_{i=1}^{n} \mathop{\mathbf{P}}_{\mathbf{v}}\mathrm{r}\left[j \notin \mathtt{SOLD}(\mathbf{v},\pi)\right] \\ \cdot \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{-i}}\left[\max\left(\mathsf{SW}_j(v_i,\widetilde{\mathbf{v}}_{-i}) - p'_j, 0\right)\cdot \mathbb{1}\left[j \in X_i(v_i,\widetilde{\mathbf{v}}_{-i})\right]\right] \\ \geq \sum_{j\in M} \mathop{\mathbf{P}}_{\mathbf{v}}\mathrm{r}\left[j \notin \mathtt{SOLD}(\mathbf{v},\pi)\right] \\ \cdot \left(\sum_{i=1}^{n} \mathop{\mathbf{E}}_{\mathbf{v}}\left[\left(\mathsf{SW}_j(\mathbf{v}) - p'_j\right)\cdot \mathbb{1}\left[j \in X_i(\mathbf{v})\right]\right]\right) \\ = \sum_{j\in M} \mathop{\mathbf{P}}_{\mathbf{v}}\mathrm{r}\left[j \notin \mathtt{SOLD}(\mathbf{v},\pi)\right]\cdot (p_j + (p_j - p'_j)). \qquad (3.3)\end{aligned}$$

In the second inequality, we decreased each probability $\mathbf{Pr}[j \notin \mathtt{SOLD}_i(\mathbf{v},\pi)]$ to $\mathbf{Pr}[j \notin \mathtt{SOLD}(\mathbf{v},\pi)]$; 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 substituted every variable $(v_i,\widetilde{\mathbf{v}}_{-i})$ to $\mathbf{v}$. The last equality follows from (3.1). Inequality (3.3) is our desired bound on the sum of buyer utilities.

We now turn to the expected revenue, which is

$$\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathrm{Rev}(\mathbf{v},\pi)\right] = \sum_{j\in M} \mathop{\mathbf{P}}_{\mathbf{v}}\mathrm{r}\left[j \in \mathtt{SOLD}(\mathbf{v},\pi)\right]\cdot (p_j - (p_j - p'_j)). \qquad (3.4)$$

Therefore, adding (3.3) and (3.4) we derive the following bound on the expected social welfare:

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

as required.

With Lemma 3.1 at hand, we are ready to complete the proof of Theorem 3.1.

Proof. [Proof Sketch] (of Theorem 3.1) It remains to show how to compute an appropriate choice of prices $\mathbf{p}'$ satisfying the conditions of Lemma 3.1. 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. In Appendix A we work out the appropriate bounds and show that $t = (\log m + \log n - \log\varepsilon)4m^2/\varepsilon^2$ samples per item are sufficient to satisfy the conditions of Theorem 3.1.

Algorithm 1 - Price computation algorithm, paramaterized by positive integer t.

For each item $j \in M$:
  Repeat t times:
    Draw $\mathbf{v} \sim \mathcal{F}$ and let $\mathbf{X} = \mathcal{A}(\mathbf{v})$.
    Let i be the agent for which $j \in X_i$.
    Query the $\mathcal{XOS}$ oracle for v_i to find $\mathsf{SW}_j(\mathbf{v})$.
  Let p'_j be half of the average value of $\mathsf{SW}_j(\mathbf{v})$ seen over all t iterations.
return $\mathbf{p}'$

We can therefore take p'_j to be the empirical estimate after this number of samples, satisfying the conditions of the theorem. To summarize, this procedure for computing $\mathbf{p}'$ is listed formally as Algorithm 1.

4 Posted Prices for General Valuations

A result similar to Theorem 3.1 holds for the more general class of $\mathcal{MPH}$-k valuations, where we get $O(k)$-approximate DSIC mechanisms for functions with complementarity level k. We will begin by formally defining

the *maximum over positive hypergraphs* ($\mathcal{MPH}$) hierarchy and providing other preliminaries. We will then provide a formal result statement.

4.1 Preliminaries and Definitions. To explain *maximum over positive hypergraphs* ($\mathcal{MPH}$) hierarchy, we first need a few preliminaries. A hypergraph representation h of valuation function $v : 2^M \to \mathbb{R}^+$ is a set function that satisfies $v(S) = \sum_{T \subseteq S} h(T)$. It is easy to verify that any valuation function v admits a unique hypergraph representation and vice versa. A set S such that $h(S) \neq 0$ is said to be a *hyperedge* of h. 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 $h(e)$. 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 h is the cardinality k of the largest hyperedge. The rank of v is the rank of its corresponding h 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 $S \subseteq M$, $h(S) \geq 0$, then we refer to function v as a *positive hypergraph-k* function (PH-k) [1]. We are now ready to present the class of $\mathcal{MPH}$-k valuations.

DEFINITION 4.1. ($\mathcal{MPH}$-k VALUATION) *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 $S \subseteq M$,*

$$v(S) = \max_{\ell \in \mathcal{L}} v_\ell(S), \tag{4.5}$$

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

It can be easily verified that 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 MPH-k oracle, which is an extension of XOS oracles to higher levels of the $\mathcal{MPH}$ hierarchy. Suppose that valuation function v is MPH-k , with supporting PH-k functions $\{v_\ell\}_{\ell \in \mathcal{L}}$. An MPH-$k$ -oracle for v takes as input a set of items S, and returns the PH-k function v_ℓ for which $v(S) = v_\ell(S)$. We will assume that this function v_ℓ 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_ℓ, 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 4.1. Note that if we restricted 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.

4.2 Pricing for $\mathcal{MPH}$-k valuations. Our result is cast in the following theorem, whose proof is deferred to Appendix B.[7]

THEOREM 4.1. *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[8] $\text{POLY}(m^k, n, 1/\varepsilon)$ *that generate expected welfare of at least $\frac{1}{4k}\mathbb{E}_{\mathbf{v} \sim \mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] - \varepsilon$ for any buyers' arrival order.*

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 C.

5 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

[7] The proof of Theorem 4.1 follows Theorem 3.1, but with an important difference: the accounting of the contribution of an item j to the welfare is more complex, since one must consider all hyperedges in which j appears. This complicates the choice of prices, as well as the derivation of welfare bounds.

[8] 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 B.

based on this information. Our proof techniques (in Theorems 3.1 and 4.1) 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 [22] 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. It remains open whether a constant approximation for subadditive valuations can be achieved.

Second, throughout the paper we assume that items are indivisible and heterogeneous. It would be an 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.

Acknowledgments

We are grateful to Shaddin Dughmi, David Kempe and Noam Nisan for helpful discussions. 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] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce*, pages 243–251, 2007.

[11] 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.

[12] 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.

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

[14] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *EC*, pages 129–138, 2010.

[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] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research, 2005.

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

[20] 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.

[21] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient.

In *STOC*, pages 201–210, 2013.

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

[23] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *J. of Economic Theory*, 87(1):95–124, 1999.

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

[25] Jason D. Hartline and Brendan Lucier. Bayesian algorithmic mechanism design. In *STOC*, 2010.

[26] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *EC*, pages 225–234, 2009.

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

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

[29] Jr Kelso, Alexander S and Vincent P Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, November 1982.

[30] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 123–136, 2012.

[31] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270 – 296, 2006. Mini Special Issue: Electronic Market Design¡/ce:title¿.

[32] Azevedo Eduardo M., Glen E. Weyl, and White Alexander. Walrasian equilibrium in large, quasi-linear markets. *Theoretical Economics*, 8:281–290, 2013.

[33] Paul Milgrom. Auctions and Bidding: A Primer. *Journal of Economic Perspectives*, 3(3):3–22, 1989.

[34] Paul Milgrom. Game theory and the spectrum auctions. *European Economic Review*, 42(3):771–778, 1998.

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

[36] Hal R. Varian. Search engine economics. In *Wirtschaftsinformatik (1)*, pages 29–30, 2007.

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

APPENDIX

A Details of the proof of Theorem 3.1

We now present the details omitted from the proof of Theorem 3.1.

We wish to show how to compute an appropriate choice of prices $\mathbf{p}'$ satisfying the conditions of Lemma 3.1. 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. We can therefore take p'_j to be the empirical estimate after this number of samples, satisfying the conditions of the theorem. This procedure is listed formally as Algorithm 1 in Section 3.

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 have that $|p'_j - p_j| < \varepsilon/2m$ for all j with probability at least $1 - \varepsilon/n$, as desired.

Setting $\delta = \varepsilon/2m$ in Lemma 3.1, 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 $\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] \le \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\sum_{i=1}^{n} v_i(M)] \le n$.

B Proof of Theorem 4.1: $\mathcal{MPH}$ valuations

We closely follow the proof of Theorem 3.1 for XOS buyers. However, there is an extra difficulty for $\mathcal{MPH}$-k valuations, since the concept of the "contribution of an item to welfare" is not as straightforward as for 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 $v_i(\cdot)$ we take the respective hypergraph representative function $A_i(\cdot)$ for the set $X_i(\mathbf{v})$, i.e., $v_i(S) \ge A_i(S)$ for any $S \subset [m]$ and $v_i(X_i(\mathbf{v})) = A_i(X_i(\mathbf{v}))$. Write $w_i(\cdot)$ for the hypergraph weights corresponding to the

hypergraph function $A_i(\cdot)$; 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 α is a constant to be determined later. The price vector $\mathbf{p}(\mathbf{v})$ has a natural interpretation: for each hyperedge in the hypergraph function $A_i(X_i(\mathbf{v}))$, divide its weight uniformly among the items in that edge; the price of item j is then the total weight allocated to item j, scaled down by factor α. 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 B.1. *For any $\mathbf{v}$, any buyer i, and any $Q \subseteq X_i(\mathbf{v})$,*

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

Proof.

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

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.

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

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

Proof. Given the prices $\mathbf{p}'$, let π 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 $\widetilde{\mathbf{v}}_{\text{-}i} \sim \mathcal{F}_{\text{-}i}$, drawn independently of $\mathbf{v}$. Consider also $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$, the allocation returned by $\mathcal{A}$ on the valuation profile $(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$. Let $S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i}) := X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) \cap \texttt{SOLD}_i(\mathbf{v}, \pi)$ 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 $R_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i}) := X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) \setminus \texttt{SOLD}_i(\mathbf{v}, \pi)$ be the subset of items in $X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})$ that remain unsold 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

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

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

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

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

$$\begin{aligned}
&\mathop{\mathbf{E}}_{\mathbf{v}} \left[\sum_i u_i(\mathbf{v}) \right] \\
&\geq \sum_i \mathop{\mathbf{E}}_{\mathbf{v}, \widetilde{\mathbf{v}}_{\text{-}i}} \left[v_i(X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})) - \sum_{j \in X_i(v_i, \widetilde{\mathbf{v}}_{\text{-}i})} p'_j \right] \\
&\quad - \alpha k \cdot \sum_i \mathop{\mathbf{E}}_{\mathbf{v}, \widetilde{\mathbf{v}}_{\text{-}i}} \left[\sum_{j \in S_i(v_i, \mathbf{v}_{\text{-}i}, \widetilde{\mathbf{v}}_{\text{-}i})} p_j(v_i, \widetilde{\mathbf{v}}_{\text{-}i}) \right].
\end{aligned} \tag{B.1}$$

Let us analyze separately the two summations on the RHS of (B.1). 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 and then using linearity of expectation, we have

$$\begin{aligned}
&\sum_i \mathop{\mathbf{E}}_{\mathbf{v},\widetilde{\mathbf{v}}_{\text{-}i}}\left[v_i(X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})) - \sum_{j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})} p'_j\right] \\
&\quad= \sum_i \mathop{\mathbf{E}}_{\mathbf{v}}\left[v_i(X_i(\mathbf{v})) - \sum_{j\in X_i(\mathbf{v})} p'_j\right] \\
&\quad= \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \sum_j p'_j \\
&\quad\geq \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \sum_j p_j - \delta m \\
&\quad= \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] \\
&\qquad - \frac{1}{\alpha}\cdot\sum_j \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i \sum_{\substack{T\ni j\\ T\subseteq X_i(\mathbf{v})}} \frac{w_i(T)}{|T|}\right] - \delta m \\
&\quad= \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \frac{1}{\alpha}\cdot\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \delta m \\
&\quad= \left(1-\frac{1}{\alpha}\right)\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] - \delta m.
\end{aligned} \tag{B.2}$$

Note that the inequality follows from the fact that $|p_j - p'_j| < \delta$ for each item j.

For the second summation on the RHS of (B.1), 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

$$\begin{aligned}
&\sum_i \mathop{\mathbf{E}}_{\mathbf{v},\widetilde{\mathbf{v}}_{\text{-}i}}\left[\sum_{j\in S_i(\mathbf{v},\widetilde{\mathbf{v}}_{\text{-}i})} p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right] \\
&\quad= \sum_i \mathop{\mathbf{E}}_{\mathbf{v},\widetilde{\mathbf{v}}_{\text{-}i}}\left[\sum_j p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right. \\
&\qquad\left. \cdot\,\mathbb{1}\left[j\in \texttt{SOLD}_i(\mathbf{v}_{\text{-}i},\pi)\right]\cdot\mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\right] \\
&\quad= \sum_{i,j} \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\cdot p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right] \\
&\qquad\cdot\mathop{\mathbf{Pr}}_{\mathbf{v}_{\text{-}i}}\left[j\in \texttt{SOLD}_i(\mathbf{v}_{\text{-}i},\pi)\right] \\
&\quad\leq \sum_{i,j} \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\cdot p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right] \\
&\qquad\cdot\mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right] \\
&\quad= \sum_j \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right] \\
&\qquad\cdot\sum_i \mathop{\mathbf{E}}_{v_i,\widetilde{\mathbf{v}}_{\text{-}i}}\left[\mathbb{1}\left[j\in X_i(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right]\cdot p_j(v_i,\widetilde{\mathbf{v}}_{\text{-}i})\right] \\
&\quad= \sum_j \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right]\cdot p_j \\
&\quad\leq \sum_j \mathop{\mathbf{Pr}}_{\mathbf{v}}\left[j\in \texttt{SOLD}(\mathbf{v},\pi)\right]\cdot p'_j + \delta m \\
&\quad= \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathrm{Rev}(\mathbf{v},\pi)\right] + \delta m.
\end{aligned} \tag{B.3}$$

The first inequality follows because the probability that j is sold before agent i arrives is dominated by the probability that j is sold at all, and the second inequality follows from the fact that $|p_j - p'_j| < \delta$ for each item j.

Substituting (B.2) and (B.3) into (B.1), we have

$$\begin{aligned}
\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right] &\geq \left(1-\frac{1}{\alpha}\right)\mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_i v_i(X_i(\mathbf{v}))\right] \\
&\quad - \delta m - \alpha k\cdot\mathop{\mathbf{E}}_{\mathbf{v}\sim\mathcal{F}}\left[\mathrm{Rev}(\mathbf{v},\pi)\right] - \alpha k\delta m.
\end{aligned} \tag{B.4}$$

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

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

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 4.1. Following the same analysis as in Theorem 3.1 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 B.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 $\mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\mathsf{SW}(\mathcal{A}(\mathbf{v}))] \le \mathbb{E}_{\mathbf{v}\sim\mathcal{F}}[\sum_{i=1}^{n} v_i(M)] \le n$.

C Lower bound for $\mathcal{MPH}$

Example. 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 $m - 1$ for the set of all m items and 0 value for any smaller subset. The optimal social welfare OPT is $m - 1$, where the second buyer is allocated all m items.

Let the seller fix prices on the items. We let the first buyer arrive first. He will buy the cheapest item, if its price is below 1. Then the second buyer has 0 value for the remaining items, which results in a social welfare of 1. In the case where each item costs at least 1, the first buyer purchases nothing but so does the second buyer, as he derives value $m - 1$ from the entire set, for a total cost of at least m. Therefore, the social welfare in the latter case is 0. We conclude that the social welfare does not exceed 1 in either of the cases, which gives us the claimed linear gap of $m - 1$ with respect to the optimal social welfare.