Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence

A Unifying Hierarchy of Valuations with Complements and Substitutes

Uriel Feige
Weizmann Institute of Science
uriel.feige@weizmann.ac.il

Michal Feldman
Tel-Aviv University
mfeldman@tau.ac.il

Nicole Immorlica
Microsoft Research
nicimm@gmail.com

Rani Izsak
Weizmann Institute of Science
ran.izsak@weizmann.ac.il

Brendan Lucier
Microsoft Research
brlucier@microsoft.com

Vasilis Syrgkanis
Microsoft Research
vasy@microsoft.com

Abstract

We introduce a new hierarchy over monotone set functions, that we refer to as $\mathcal{MPH}$ (Maximum over Positive Hypergraphs). Levels of the hierarchy correspond to the degree of complementarity in a given function. The highest level of the hierarchy, $\mathcal{MPH}$-m (where m is the total number of items) captures all monotone functions. The lowest level, $\mathcal{MPH}$-1, captures all monotone submodular functions, and more generally, the class of functions known as $\mathcal{XOS}$. Every monotone function that has a positive hypergraph representation of rank k (in the sense defined by Abraham, Babaioff, Dughmi and Roughgarden [EC 2012]) is in $\mathcal{MPH}$-k. Every monotone function that has supermodular degree k (in the sense defined by Feige and Izsak [ITCS 2013]) is in $\mathcal{MPH}$-$(k+1)$. In both cases, the converse direction does not hold, even in an approximate sense. We present additional results that demonstrate the expressiveness power of $\mathcal{MPH}$-k.

One can obtain good approximation ratios for some natural optimization problems, provided that functions are required to lie in low levels of the $\mathcal{MPH}$ hierarchy. We present two such applications. One shows that the maximum welfare problem can be approximated within a ratio of $k+1$ if all players hold valuation functions in $\mathcal{MPH}$-k. The other is an upper bound of $2k$ on the price of anarchy of simultaneous first price auctions.

1 Introduction

In a combinatorial auction setting, a set M of m items is to be allocated among a set N of n buyers. Each buyer $i \in N$ has a valuation function that assigns a non-negative real number $v_i(S)$ to every bundle of items $S \subseteq M$. A well motivated objective is to find a partition of the items $X = (X_1, \ldots, X_n)$ among the buyers so as to maximize the *social welfare*, defined as the sum of buyers' valuations from the bundles they obtain $SW(X) = \sum_{i \in N} v_i(X_i)$. The model of combinatorial auctions is highly applicable to real-world settings such as spectrum auctions and electronic advertisement markets.

Most of the existing literature on combinatorial auctions has focused on the case where buyer valuations are complement-free. Roughly speaking, this means that the value for the union of two bundles of items cannot exceed the sum of the values for each individual bundle. Such valuations do not capture scenarios where certain items produce more value when acquired in conjunction with each other (such as a left and right shoe). Complement-free valuations are arguably more well-behaved than general combinatorial valuations in many aspects. From an algorithmic perspective, complement-free valuations admit constant-factor polynomial time approximation algorithms (Feige and Vondrák 2006; Dobzinski and Schapira 2006; Feige 2006), while general valuations are hard to approximate even to within a factor that is sub-polynomial in the number of items (Lehmann, O'Callaghan, and Shoham 1999). From a game-theoretic perspective, simple auctions, such as running simultaneously a single-item first-price auction for each item, induce equilibria that achieve constant factor approximations to the optimal welfare if all valuations are complement-free (Syrgkanis and Tardos 2013; Feldman et al. 2013; Hassidim et al. 2011; Christodoulou, Kovács, and Schapira 2008; Bhawalkar and Roughgarden 2011). In contrast, if valuations exhibit complements, the worst-case inefficiency grows with the number of items (Hassidim et al. 2011).

While the theory suggests that complementarities degrade the performance of combinatorial auctions, they arise very naturally in many economic scenarios. A prominent example is the FCC spectrum auctions, where it is desirable to win licenses for the same band of spectrum in adjacent geographical regions. The prevalence of complementarities in practice calls for a better theoretical understanding of the effect of the level of complementarity on auction performance.

To this end, we introduce a new hierarchy of monotone set functions called *maximum over positive hypergraphs* ($\mathcal{MPH}$), whose level captures the degree of complementarity. A new hierarchy is useful if it has a strong expressiveness power on the one hand, and algorithmic and economic implications on the other. We show that important classes of functions are captured in low levels of our hierarchy (a detailed exposition is deferred to Section 1.2). We then present algorithmic and economic results that illustrate the usefulness of our hierarchy. In particular, we develop an algorithm that approximates the welfare maximization problem to within a factor of $k+1$, where k is the degree of complementarity. We further show that an auction that solicits bids on each item separately and allocates each item to the highest bidder (at a cost equals to her bid) achieves a $2k$-approximation to the optimal welfare at any equilibrium of bidder behavior.

1.1 The Maximum over Positive Hypergraph ($\mathcal{MPH}$) Hierarchy

Given a set M of m items, a set function $v : 2^M \to \mathbb{R}^+$ is *normalized* if $v(\emptyset) = 0$ and *monotone* if $v(T) \geq v(S)$ whenever $S \subseteq T \subseteq M$.[1] A normalized monotone set function is necessarily non-negative. Throughout the paper we assume that all set functions are normalized and monotone, unless stated otherwise. In the context of combinatorial auctions, we refer to the set functions as valuation functions.

A set function v is symmetric if $v(S) = v(T)$ whenever $|S| = |T|$. A hypergraph representation of a set function $v : 2^M \to \mathbb{R}^+$ is a (normalized but not necessarily monotone) set function $h : 2^M \to \mathbb{R}$ that satisfies $v(S) = \sum_{T \subseteq S} h(T)$. It is easy to verify that any set 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. Pictorially, the hypergraph representation can be thought of 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 largest cardinality of any hyperedge. Similarly, the *positive rank* (respectively, *negative rank*) of h is the largest cardinality of any hyperedge with strictly positive (respectively, negative) value. The rank of a set function v is the rank of its corresponding hypergraph representation, and we refer to a function v with rank r as a *hypergraph-r* function. Last, if the hypegraph representation is non-negative, i.e. for any $S \subseteq M$, $h(S) \geq 0$, then we refer to such a function as a *positive hypergraph-r* ($\mathcal{PH}$-r) function.

We define a parameterized hierarchy of set functions, with a parameter corresponding to the degree of complementarity.

Definition 1 (Maximum Over Positive Hypergraph-k ($\mathcal{MPH}$-k) class). *A monotone set 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 $\mathcal{PH}$-k functions. That is, there exist $\mathcal{PH}$-k functions $\{v_\ell\}_{\ell \in \mathcal{L}}$ such that for every set $S \subseteq M$,*

$$v(S) = \max_{\ell \in \mathcal{L}} v_\ell(S), \tag{1}$$

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

The $\mathcal{MPH}$ hierarchy has the following attributes:

1. **Completeness**. Every monotone set function is contained in some level of the hierarchy (see below).
2. **Usefulness**. The hierarchy has implications that relate the level in the hierarchy to the efficiency of solving optimization problems. Specifically, we show implications of our hierarchy to the approximation guarantee of the algorithmic welfare maximization problem (in Section 3.2) and the price of anarchy of simultaneous single item auctions (in Section 3.3).
3. **Expressiveness**. The hierarchy is expressive enough to contain many functions in its lowest levels (see Section 2).

We conclude this section with some basic properties of the $\mathcal{MPH}$ hierarchy (for more properties, see Section 2). The two extreme cases of $\mathcal{MPH}$-k functions coincide with two important classes of valuations. Specifically, $\mathcal{MPH}$-1 is the class of functions that can be expressed as the maximum over a set of additive functions. This is exactly the class of $\mathcal{XOS}$ valuations (Lehmann, Lehmann, and Nisan 2001), which is a complement-free valuation class that has been well-studied in the literature. This class contains all submodular valuations, i.e. valuations that exhibit decreasing marginal returns. On the other side, $\mathcal{MPH}$-m coincides with the class of all monotone functions,[2] and so the hierarchy is complete. For intermediate values of k, $\mathcal{MPH}$-k is monotone; namely, for every $k < k'$ it holds that $\mathcal{MPH}$-$k \subset \mathcal{MPH}$-k'. We get the following hierarchy:

$$\mathcal{XOS} = \mathcal{MPH}\text{-}1 \subset \cdots \subset \mathcal{MPH}\text{-}m = \text{Monotone} \tag{2}$$

A simple example. Consider the following example, which has an intuitive interpretation in the context of FCC spectrum auctions. There are four items for sale, $\{A_1, A_2, B_1, B_2\}$, corresponding to spectrum bands A and B in each of two neighboring geographic regions $i = 1, 2$. The value of any A_i is 1, but a bidder has value 3 for getting both A_1 and A_2, due to the complementary relationship of being in neighboring regions. The value of band B is similar to A. However, bands A and B are substitutes, so if a bidder obtains items from each band then he can derive value from only one of them. This valuation can be represented as a hypergraph with four nodes A_1, A_2, B_1, B_2, where the weight of each single node is 1, the weight of the edges $(A_,1, A_2)$ and (B_1, B_2) are 1, the weight of the edges (A_1, B_1) and (A_2, B_2) are -1, and the weight of the hyperedge including all four nodes is -1. This valuation can also be represented as a maximum over positive hypergraph valuations of rank 2, using the following four graphs: Graph G_1 assigns weight 0 to nodes B_i, weight 1 to nodes A_i and weight 1 to the edge (A_1, A_2). Graph G_2 is the same as G_1, swapping A and B. Graph G_3 assigns weight 0 to nodes A_2 and B_1, and weight 1 to nodes A_1 and B_2. Graph G_4 assigns weight 1 to nodes A_2 and B_1, and weight 0 to nodes A_1 and B_2. Consider, for example, the set $\{A_1, A_2\}$. Its value is 3 according to the hypergraph valuation, and indeed, it obtains maximum value in G_1, which assigns it value 3, as desired. This valuation is therefore $\mathcal{MPH}$-2.

Fractionally "Subadditive" Characterization of $\mathcal{MPH}$-k. In the full version of the paper, we show the definition of $\mathcal{MPH}$-k functions has a natural analogue as an extension of *fractionally subadditive* functions.

1.2 Related Work

Expressiveness. Since the maximum welfare allocation problem is $\mathcal{NP}$-hard to approximate even with very poor ratio (see for example (Lehmann, O'Callaghan, and Shoham

[1]We use $\mathbb{R}^+$ for non-negative real numbers. That is, 0 is included.

[2]Simply create a separate $\mathcal{PH}$-$|S|$ function for each set S with a single hyperedge equal to the set S and with weight $f(S)$. Then, by monotonicity, the maximum of these functions is equal to the initial valuation.

1999) for the case of *single minded bidders* – bidders that want one particular bundle of items), there has been extensive work on classification of monotone set functions. We distinguish between two types of classifications. One is that of *restricted classes* of set functions, and the other is *inclusive hierarchies* that capture all monotone set functions.

Restricted classes of monotone set functions. (Lehmann, Lehmann, and Nisan 2001) initiated a systematic classification of set functions without complementarities (e.g. *gross-substitutes* (Kelso and Crawford 1982), *submodular*, $\mathcal{XOS}$ (Sandholm 1999), *subadditive*). Subsequent research gave constant factor approximation algorithms for these classes ($1 - 1/e - \epsilon$ in the submodular case (Feige and Vondrák 2006), $1 - 1/e$ in the $\mathcal{XOS}$ case (Dobzinski and Schapira 2006) and 2 in the subadditive case (Feige 2006)). These algorithms assume *demand query* access to the valuation functions, though for the submodular case, if one is satisfied with a $1 - 1/e$ ratio, then *value queries* suffice (Vondrák 2008). The first level of our hierarchy, i.e. $\mathcal{MPH}$-1, coincides with the class of $\mathcal{XOS}$ functions.

(Conitzer, Sandholm, and Santi 2005) consider *graphical valuations*: every item has a weight, and every pair of items (edge of the graph) has a weight (positive or negative) and the value of a set of items is the sum of weights of items and edges within the set. We show that this class is in $\mathcal{MPH}$-2.

(Abraham et al. 2012) consider the hierarchy of $\mathcal{PH}$-k valuation functions, as already defined, (which are obviously contained in $\mathcal{MPH}$-k) that allows only complements but no substitutes (e.g. submodular functions cannot be expressed in this hierarchy, and even some supermodular functions cannot be expressed). They show that the maximum welfare problem can be approximated within a ratio of k if all valuation functions are in $\mathcal{PH}$-k and they design truthful mechanisms that achieve approximation ratios that degrade logarithmically with the numer of items m.

Complete hierarchies of monotone set functions. (Feige and Izsak 2013) introduced a hierarchy parameterized by the so-called *supermodular degree* (the lowest level coinciding with submodular functions) and gave a greedy $(k + 2)$-approximation algorithm for the welfare maximization problem when functions are in the k-th level of the hierarchy. We show that every level of the supermodular degree hierarchy is strictly contained in the corresponding level of the $\mathcal{MPH}$ hierarchy,whilst there are functions in $\mathcal{MPH}$-2 that cannot even be approximated by functions of low supermodular degree (e.g. functions of supermodular degree $\sqrt{m}$ approximate them only within a ratio of $\Omega(\sqrt{m})$).

The $\mathcal{XOS}$ class introduced in (Lehmann, Lehmann, and Nisan 2001) is based on "OR" and "XOR" operations previously introduced in (Sandholm 1999), but with the restriction that "OR" operations are applied on single items. Removing this restriction and allowing operations on bundles, one obtains an $\mathcal{XOS}$ hierarchy parameterized by the size of the largest bundle. While $\mathcal{XOS}$-1 coincides with $\mathcal{MPH}$-1, $\mathcal{MPH}$-k is strictly larger than $\mathcal{XOS}$-k, i.e. $\mathcal{XOS}$-k is contained in $\mathcal{MPH}$-k, while there are functions in $\mathcal{MPH}$-2 that cannot be approximated in $\mathcal{XOS}$-k for any constant k.

Welfare approximation. The complement-free valuations introduced in (Lehmann, Lehmann, and Nisan 2001) have also been studied in the game-theoretic context of equilibria in simultaneous single-item auctions. It has been established that the Bayes-Nash and Correlated price of anarchy of this auction format, with a first-price payment rule, are at most $\frac{e}{e-1}$ in the $\mathcal{XOS}$ case (Syrgkanis and Tardos 2013) and at most 2 in the subadditive case (Feldman et al. 2013). For the second-price payment rule, these bounds become 2 for $\mathcal{XOS}$ (Christodoulou, Kovács, and Schapira 2008) and 4 for subadditive (Feldman et al. 2013). These results build upon a line of work studying non-truthful item auctions for complement-free valuations (Bikhchandani 1999; Bhawalkar and Roughgarden 2011; Christodoulou, Kovács, and Schapira 2008; Hassidim et al. 2011; Paes Leme, Syrgkanis, and Tardos 2012). Equilibrium analysis of non-truthful auctions has been applied to several mechanism design settings (see (Caragiannis et al. 2012; Lucier and Borodin 2010; Markakis and Telelis 2012)) and recent papers give general frameworks for bounding the inefficiency of mechanisms at equilibrium (Roughgarden 2012; Syrgkanis 2012; Syrgkanis and Tardos 2013).

2 Summary of Results

We obtain results on the expressiveness power of the $\mathcal{MPH}$ hierarchy and show applications of it for approximating social welfare in combinatorial auctions.

Expressiveness. The first theorem establishes the expressiveness power of $\mathcal{MPH}$.

Theorem 1. *The $\mathcal{MPH}$ hierarchy captures many existing hierarchies, as follows:*

1. *By definition, $\mathcal{MPH}$-1 is equivalent to the class $\mathcal{XOS}$ (Lehmann, Lehmann, and Nisan 2001) and every function that has a positive hypergraph representation of rank k (Abraham et al. 2012) is in $\mathcal{MPH}$-k.*
2. *Every monotone graphical valuation (Conitzer, Sandholm, and Santi 2005) is in $\mathcal{MPH}$-2. Furthermore, every monotone function with positive rank 2 is in $\mathcal{MPH}$-2.*
3. *Every monotone function that has a hypergraph representation with positive rank k and laminar negative hyperedges (with arbitrary rank) is in $\mathcal{MPH}$-k.*
4. *Every monotone function that has supermodular degree k (Feige and Izsak 2013) is in $\mathcal{MPH}$-$(k+1)$.*

We further establish that the converse direction does not hold, even in an approximate sense, and conclude that the $\mathcal{MPH}$-k hierarchy is *strictly* more expressive than many existing hierarchies. Specifically, we show that $\mathcal{MPH}$-1 and $\mathcal{MPH}$-2 contain functions that cannot be *approximated* by functions in low levels of other hierarchies.

Definition 2. *We say that a set function f approximates a set function g within a ratio of $\rho \geq 1$ if there are ρ_1 and ρ_2 such that for every set S $\rho_1 \leq \frac{f(S)}{g(S)} \leq \rho_2$, and $\frac{\rho_2}{\rho_1} \leq \rho$.*

Proposition 2. *There are functions in very low levels of the $\mathcal{MPH}$ hierarchy that cannot be approximated well even at relatively high levels of other hierarchies, as follows:*

1. *There exists a submodular function (i.e., supermodular degree 0, $\mathcal{MPH}$-1) such that*

(a) A graphical function cannot approximate it within a ratio better than $\Omega(m)$.

(b) A positive hypergraph function cannot approximate it within a ratio better than m.

(c) A hypergraph function of rank k (both negative and positive) cannot approximate it within a ratio better than $\Omega(\frac{m}{k^2})$, for every k.

2. There exists a $\mathcal{PH}$-2 function (i.e., $\mathcal{MPH}$-2) such that every function of supermodular degree d cannot approximate it within a ratio better than $\Omega(m/d)$.

Applications. With the new hierarchy at hand, we are in a position to revisit algorithmic and game-theoretic problems about welfare maximization in combinatorial auctions. We obtain good approximation ratios for settings with valuations that lie in low levels of the $\mathcal{MPH}$-k hierarchy. We first provide a polynomial time algorithm for the welfare maximization problem when valuations are in $\mathcal{MPH}$-k.

Theorem 3. *If agents have $\mathcal{MPH}$-k valuations, then there exists an algorithm that gives $k+1$ approximation to the optimal social welfare. This algorithm runs in polynomial time given an access to demand oracles for the valuations.*

Our approximation algorithm first solves the configuration linear program for welfare maximization introduced by (Dobzinski, Nisan, and Schapira 2010). Solving this LP can be done in polynomial time using demand queries. We then round the solution to the LP so as to get an integer solution. Our rounding technique is oblivious and does not require access to demand queries. By analyzing the integrality gap, we show that our rounding technique is nearly best possible.

The second setting we consider is a simultaneous first-price auction — where each of the m items is sold via a separate single-item auction. We quantify the welfare loss in this simple auction when bidders have $\mathcal{MPH}$-k valuations and at every coarse correlated equilibrium of the complete information setting (correlated price of anarchy) and Bayes-Nash equilibrium of the incomplete information setting (Bayes-Nash price of anarchy).

Theorem 4. *For simultaneous first price auctions with $\mathcal{MPH}$-k valuations, both the correlated price of anarchy and the Bayes-Nash price of anarchy are at most $2k$.*

Our proof technique extends the analysis for complement-free valuations in (Feldman et al. 2013) and the smoothness framework introduced in (Syrgkanis and Tardos 2013) to settings with complementarities. We also establish an almost matching lower bound in the full version of the paper.

Theorem 5. *There exists an instance of a simultaneous first price auction with single minded bidders in $\mathcal{MPH}$-k in which the price of anarchy is $\Omega(k)$.*

Remarks. Most of our expressiveness results showing that a certain function belongs to $\mathcal{MPH}$-k are established by showing that the function satisfies a certain requirement that we refer to as the Positive Lower Envelope (PLE) condition. We also observe that, together with monotonicity, this requirement becomes a sufficient and necessary condition for membership in $\mathcal{MPH}$-k. This observation motivates the definition of a new hierarchy, referred to as $\mathcal{PLE}$. The class $\mathcal{PLE}$-k contains $\mathcal{MPH}$-k, but also includes non-monotone functions. While monotonicity is a standard assumption in the context of combinatorial auctions, $\mathcal{PLE}$ can be applicable outside the scope of combinatorial auctions. In the full version of the paper, we analyze the expressiveness of $\mathcal{PLE}$ functions and the observation that our approximation results extend to non-monotone $\mathcal{PLE}$ functions.

Extensions. One of the main open problems suggested by this work is the relation between hypergraph valuations of rank k and $\mathcal{MPH}$-k valuations. We conjecture that:

Conjecture 6. *Every hypergraph function with rank k (positive or negative) is in $\mathcal{MPH}$-$O(k^2)$.*

We make partial progress toward the proof of this conjecture, by confirming it for the case of symmetric functions. For non-symmetric, observe that for the case of laminar negative hyperedges, we show an even stronger statement in item (3) of Theorem 1.

Theorem 7. *Every monotone symmetric hypergraph function with rank k (positive or negative) is in $\mathcal{MPH}$-$O(k^2)$*

For symmetric functions, we conjecture a more precise bound of $\lceil \frac{k}{2} \rceil \lceil \frac{k+1}{2} \rceil$, suggested by a computer-aided simulation based on a non-trivial LP formulation. For the special cases of symmetric functions of ranks $k = 3$ and 4, we show that they are in $\mathcal{MPH}$-4 and $\mathcal{MPH}$-6, respectively, and that this is tight. We use an LP formulation whose optimal solution is the worst symmetric function possible for a given rank, and its value corresponds to the level of this worst function in the $\mathcal{MPH}$ hierarchy. We bound the value of this LP, by using LP duality (in the full version).

3 Proofs

In this section we include a part of our proofs. Due to space constraints, we defer the other proofs to the appendix.

3.1 Some proofs of expressiveness (Thm. 1 and 7)

Positive lower envelope technique. Proving that a particular set function $f : 2^M \to \mathbb{R}^+$ is in $\mathcal{MPH}$-k requires constructing a set of $\mathcal{PH}$-k valuations that constitute the index set $\mathcal{L}$ over which the maximum is taken. In what follows we present a canonical way of constructing the set $\mathcal{L}$. The idea is to create a $\mathcal{PH}$-k function for every subset S of the ground set M. The collection of these $\mathcal{PH}$-k functions, one for each subset, constitutes a valid $\mathcal{MPH}$-k representation if they adhere to the following condition.

Definition 3 (Positive Lower Envelope (PLE)). *Let $f : 2^M \to \mathbb{R}^+$ be a monotone set function. A positive lower envelope (PLE) of f is any positive hypergraph function g such that, $g(M) = f(M)$ and for any $S \subseteq M$: $g(S) \leq f(S)$ (no-overestimate).*

Before presenting the characterization, we need the following definition. A function $f : 2^M \to \mathbb{R}^+$ *restricted to a subset* S, $S \subseteq M$, is a function $f_S : 2^S \subseteq \mathbb{R}^+$ with $f_S(S') = f(S')$ for every $S' \subseteq S$.

Proposition 8 (A characterization of $\mathcal{MPH}$). *A function f is in $\mathcal{MPH}$-k if and only if it is monotone and f_S admits a lower envelope of rank k for every set $S \subseteq M$.*

We provide a proof sketch of the second part in Thm. 1, that any monotone function of positive rank 2 is in $\mathcal{MPH}$-2

Proof. Let $v : 2^M \to \mathbb{R}^+$ be a monotone set function of positive rank 2 and let G_v be the hypergraph representation of v, where the vertices of G_v are the items of M. By Proposition 8 it suffices to show that every $S \subseteq M$ has a positive lower envelope of rank 2 (abbreviated as PLE-2). Consider an arbitrary $S \subseteq M$. We construct a positive lower envelope for S by induction. Starting with an empty set of vertices, we iteratively add the vertices of S, one at a time. Let $u_i \in S$ denote the vertex added at iteration i, and $S_i \subseteq S$ denote the resulting subset. The inductive invariant that we maintain is that each S_i has a PLE-2. The base case of the induction is S_1, and there the inductive hypothesis holds because v is nonnegative. We now prove the inductive step. Namely, we assume that S_{i-1} has a PLE-2, and prove the same for S_i.

Let N_i (P_i, respectively) denote the set of negative (positive, respectively) hyperedges in G_v that contain u_i and are contained in S_i. (As v has positive rank 2, the hyperedges in P_i have rank at most 2.) Consider an auxiliary bipartite graph H with members of N_i as one set of vertices, members of P_i as the other set of vertices, and edges between $e \in N_i$ and $e' \in P_i$ iff $e' \subset e$ (namely, the negative hyperedge contains all items of the positive hyperedge). These edges have infinite capacities. Add two auxiliary vertices, s connected to each member of N_i by an edge of capacity equal to the (absolute value of the) weight of the corresponding hyperedge in G_v, and t connected to each member of P_i by an edge of capacity equal to the weight of the corresponding hyperedge in G_v. We claim that there is a flow F from s to t saturating all edges of s. This follows from the max flow min cut theorem, together with the facts that v is monotone and all positive hyperedges have rank at most 2. Given this claim (whose proof appears in the full version), we add to the PLE-2 of S_{i-1} only the the members of P_i (hence positive edges of rank at most 2), but each of them with a weight reduced by the amount of flow that goes from it to t (according to the saturated flow F). The flow F gives us a way of charging every negative hyperedge that is discarded against a reduction in weight of positive hyperedges contained in it, and this implies that the result is indeed a PLE-2 for S_i. ∎

Next we provide a proof sketch of Thm. 7, that any monotone symmetric hypergraph-r function is in $\mathcal{MPH}$-$O(r^2)$.

Proof. Let f be a normalized monotone symmetric set function of rank r, and let h be its hypergraph representation. Consider the following monotone symmetric set function g defined by the positive hypergraph representation p: $p(S) = f(U)/\binom{n}{R}$ if $|S| = R$, and $p(S) = 0$ o.w., for $R = 3r^2$. As all four functions f, h, g, p are symmetric, we shall change notation and replace $f(S)$ by $f(|S|)$. As special cases of this notation, $f(U)$ is replaced by $f(n)$, and $f(\phi)$ by $f(0)$.

We claim that g is a lower envelope for f. There are three conditions to check. Two of them trivially hold, namely, $g(0) = f(0) = 0$, and $g(n) = \binom{n}{R}p(R) = f(n)$. The remaining condition requires that $g(k) \leq f(k)$ for every $1 \leq k \leq n-1$. This holds for $k < R$, since $g(k) = 0$, whereas $f(k) \geq 0$. Hence the main content of our proof is to establish the inequality $g(k) \leq f(k)$ for every $R \leq k \leq n-1$.

The proof proceeds by means of contradiction: suppose there is some f that serves as a negative example, namely, that for this f there is $R \leq k \leq n-1$ for which $g(k) > f(k)$. We can show that if such an example exists then there exists one where $k = n-1$ (details appear in the full version). Thus it suffices to show that $g(n-1) = \binom{n-1}{R}f(n)/\binom{n}{R} = \frac{n-R}{n}f(n) \leq f(n-1)$ for any f that is hypergraph-r.

We will consider the (not necessarily monotone) degree r polynomial $F(x) = \sum_{i=1}^{r}\binom{x}{i}h(i)$, that matches $f(x)$ at integral points $\{0,\ldots,n\}$. Let $M = \max_{0\leq x\leq n}|F(x)|$ and let $0 \leq y \leq n$ be such that $|F(y)| = M$. By Markov's inequality regarding bounds on derivatives of polynomials (Markov 1890), we can show that $\max_{0\leq x\leq n}|F'(x)| \leq \frac{2r^2}{n}M$. If y is an integer then monotonicity of f (and hence of F on integer points) implies that $M = f(n)$. However, y need not be integer. In that case $i < y < i+1$ for some $0 \leq i \leq n-1$. Let $m = \max[|F(i)|, |F(i+1)|]$. Then $M \leq m + \frac{1}{2}\max_{i\leq x\leq i+1}[|F'(x)|] \leq f(n) + \frac{r^2}{n}M$. As $n \geq R \geq 3r^2$ we obtain that $M \leq 3f(n)/2$. On the other hand, $f(n-1) = F(n-1) \geq f(n) - \max_{0\leq x\leq n}F'(x) \geq f(n) - \frac{2r^2}{n}M \geq f(n) - \frac{3r^2}{n}f(n)$. Since $R = 3r^2$ we have that $f(n-1) \geq (1-\frac{R}{n})f(n) = g(n-1)$, as desired. ∎

3.2 Algorithmic Welfare Maximization (Thm. 3)

In this section we consider the purely algorithmic problem, ignoring incentive constraints. While constant factor approximations exist for welfare maximization in the absence of complementarities (see (Dobzinski, Nisan, and Schapira 2010; Feige 2006)), it is not hard to see that complementarities can make the welfare problem as hard as independent set and hence inapproximable to within an almost linear factor. Our hierarchy offers a linear degradation of the approximation as a function of the degree of complementarity. At a high level, our algorithm works as follows: define the *configuration linear program* (LP) (introduced in (Dobzinski, Nisan, and Schapira 2010)) by introducing a variable $x_{i,S}$ for every agent i and subset of items S. Given the valuation function v_i of each agent i, the *configuration LP* is:

$$\begin{aligned} \text{maximize} \quad & \textstyle\sum_{i,S} x_{i,S}\cdot v_i(S) & (3)\\ \text{s.t.} \quad & \textstyle\sum_{S} x_{i,S} \leq 1 \quad \forall i \in N \\ & \textstyle\sum_{i,S|j\in S} x_{i,S} \leq 1 \quad \forall j \in M \\ & x_{i,S} \geq 0 \quad \forall i \in N, S \subseteq M \end{aligned}$$

The first set of constraints guarantees that no agent is allocated more than one set and the second set of constraints guarantees that no item belongs to more than one set. This LP provides an upper bound on the optimal welfare. To find a solution that approximates the optimal welfare, we first solve this LP (through duality using *demand queries* (Nisan and Segal 2006)) and then round it (see below).

Rounding the LP. First each agent i is assigned a tentative set S_i' according to the probability distribution induced by the variables $x_{i,S}$. Note that this tentative allocation has the same expected welfare as the LP. However, it may be infeasible as agents' sets might overlap. We must resolve these contentions. Several approaches for doing this when there are no complementarities were proposed and analyzed in (Dobzinski, Nisan, and Schapira 2010; Feige 2006). However, these approaches will fail badly in our setting, due to the existence of complementarities. Instead, we resolve contention using the following technique: We generate a uniformly random permutation π over the agents and then at each step t for $1 \leq t \leq n$, assign agent $i = \pi(t)$ items $S_i = S_i' \setminus \{\cup_{i'=\pi(1)}^{\pi(t-1)} S_{i'}\}$, i.e., those items in his tentative set that have not already been allocated.

The following proposition bounds the welfare guarantee of the above contention resolution algorithm.

Proposition 9. *If agents have $\mathcal{MPH}$-k valuations, then the random permutation rounding algorithm produces (in expectation) an allocation that approximates the maximum welfare within a ratio no worse than $k+1$.*

Proof. First, note that the solution is feasible, since every item is allocated at most once. We upper bound the approximation guarantee. The sum of values of tentative sets preserve, in expectation, the value of the optimal welfare returned by the configuration LP. Consider an arbitrary agent and his tentative set T. This set attained its value according to some positive hypergraph H that has no edges of rank larger than k. Consider an arbitrary edge of H contained in T, and let $r \leq k$ be its rank. We claim that its expected contribution (expectation taken over the random choices of the other agents and the random permutation) towards the final welfare is at least $1/(r+1)$ of its value. The expected number of other agents who compete on items from this edge is at most r (by summing up the fractional values of sets that contain items from this edge). Given that there are ℓ other competing agents, the agent gets all items from the edge with probability exactly $\frac{1}{\ell+1}$. As the expectation of ℓ is at most r, the expectation of $\frac{1}{\ell+1}$ is at least $\frac{1}{r+1}$ (by convexity) and hence at least $\frac{1}{k+1}$ as the valuation function is $\mathcal{MPH}$-k. The proof follows from linearity of expectation. ∎

An integrality gap of $k - 1 + \frac{1}{k}$ is known for hypergraph matching in k-uniform hypergraphs (see (Chan. Y. H 2012) and references therein). These instances are special cases of welfare maximization with $\mathcal{MPH}$-k valuations. Hence, our rounding technique in Proposition 9 is nearly best possible. Recall also that even for single-minded bidders with sets of size up to k, it is $\mathcal{NP}$-hard to approximate the welfare maximization problem to a better factor than $\Omega(\frac{\ln k}{k})$.[3]

3.3 Welfare at Equilibrium (Thm. 4)

In this section we study welfare guarantees at equilibrium of the simultaneous item auction, when all agents have $\mathcal{MPH}$-k valuations. In a simultaneous item (first-price) auction, every bidder $i \in [n]$ simultaneously submits a bid $b_{ij} \geq 0$ for every item $j \in [m]$. We write $b_i = (b_{i1}, \ldots, b_{im})$ for the vector of bids of bidder i, and $b = (b_1, \ldots, b_n)$ for the bid profile of all bidders. Every item is allocated to the bidder who submits the highest bid on it (breaking ties arbitrarily), and the winning bidder pays his bid. We let $X_i(b)$ denote the bundle allocated to bidder i under bid profile b, and we write $X(b) = (X_1(b), \ldots, X_n(b))$ for the allocation vector under bids b. When clear in the context, we omit b and write X for the allocation. A bidders' utility is assumed to be quasilinear; i.e., $u_i(b; v_i) = v_i(X_i(b)) - \sum_{j \in X_i(b)} b_{ij}$. Given a valuation profile $\mathbf{v}$ we let X^* be the welfare-maximizing allocation and $\textsc{Opt}(\mathbf{v})$ its social welfare.

In this part we assume that the valuations of the players are common knowledge and we provide a weaker efficiency guarantee. We defer the tighter analysis and the extension to the incomplete information setting to the full version. A Nash equilibrium is a profile of (possibly random) bids $B = (B_1, \ldots, B_n)$, such that no player's utility can increase by deviating to another bid. To quantify the inefficiency of a simultaneous item auction, we will use the *price of anarchy* (PoA) measure, which is the maximum ratio (over all valuation profiles) of the optimal welfare over the welfare obtained at any Nash equilibrium.

$$\text{PoA} = \max_{\mathbf{v},\ B:\ B \text{ is mixed NE}} \frac{\textsc{Opt}(\mathbf{v})}{\mathbb{E}_{\mathbf{b}\sim B}[SW(X(\mathbf{b}))]}. \quad (4)$$

Bounding the PoA. We provide a proof that the PoA of the auction is at most $4k$, when bidders have $\mathcal{MPH}$-k valuations. Let B be a randomized bid profile that constitutes a Nash equilibrium under valuations $\mathbf{v}$. For each item $j \in [m]$, let $P_j = \max_j B_{ij}$ be the price of item j; P_j is a random variable induced by the bid profile. Consider what would happen if bidder i deviated from B and instead bid $b^*_{ij} = 2k \cdot \mathbb{E}[P_j]$ on all the items $j \in X^*_i$ and 0 on the other items. By Markov's inequality bidder i wins each item $j \in X^*_i$ with probability at least $1 - \frac{1}{2k}$. Let v^*_i be the $\mathcal{PH}$-k lower envelope with respect to set X^*_i (recall bidders have $\mathcal{MPH}$-k valuations). Then, $v_i(X^*_i) = v^*_i(X^*_i)$ and, for any $X_i \subseteq X^*_i$, $v_i(X_i) \geq v^*_i(X_i)$. Since v^*_i is a $\mathcal{PH}$-k valuation, each hyperedge of v^*_i has size at most k; it then follows by the union bound that bidder i wins all items in any such hyperedge with probability at least $\frac{1}{2}$. Therefore, the value that the player derives from this deviation is at least $\frac{1}{2}v^*_i(X^*_i) = \frac{1}{2}v_i(X^*_i)$. Hence, his utility from the deviation is at least $\frac{1}{2}v_i(X^*_i) - 2k \cdot \sum_{j \in X^*_i} \mathbb{E}[P_j]$. By the Nash condition his utility at equilibrium is at least this high.

Summing the above bound over all bidders i, the sum of bidders' utilities at equilibrium is at least $\frac{1}{2}\textsc{Opt}(v) - 2k \cdot \sum_{j \in [m]} \mathbb{E}[P_j]$. Since total utility is welfare minus revenue:

$$\mathbb{E}[SW(B; \mathbf{v})] - \sum_{j \in [m]} \mathbb{E}[P_j] \geq \frac{1}{2}\textsc{Opt}(v) - 2k \sum_{j \in [m]} \mathbb{E}[P_j]$$

Since every player has the option to drop out of the auction, his expected utility is non-negative. Therefore, the expected total payment at equilibrium is bounded above by the

[3]Follows by an approximation preserving reduction from k-set packing from (Lehmann, O'Callaghan, and Shoham 1999), together with a hardness result of (Hazan, Safra, and Schwartz 2006).

welfare.Substituting this in the above inequality gives that $2k \cdot \mathbb{E}[SW(B; \mathbf{v})] \geq \frac{1}{2}\textsc{Opt}(v)$, which establishes an upper bound of $4k$ on the PoA, as desired.

Acknowledgments

We are grateful to Noam Nisan for valuable discussions. We used computer-assisted simulations in order to find feasible solutions for the dual LPs. Specifically, we used Microsoft .NET together with Gurobi (Gurobi Optimization 2014) and IBM CPLEX and also Wolfram Mathematica. We thank Ofer Bartal and Matan Karklinsky for their help in this matter. The work of Uriel Feige and Rani Izsak is supported in part by the Israel Science Foundation (grant No. 621/12) and by the I-CORE Program of the Planning and Budgeting Committee and The Israel Science Foundation (grant No. 4/11). The work of Rani Izsak is also supported in part by Irit Dinur's ERC-StG grant number 239985. The work of Michal Feldman was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122. The work of Vasilis Syrgkanis is partly supported by a Simons Graduate Fellowship in Theoretical Computer Science and part was done while an intern at Microsoft Research. The work of Nicole Immorlica was partly supported by the Alfred P. Sloan Foundation.

References

Abraham, I.; Babaioff, M.; Dughmi, S.; and Roughgarden, T. 2012. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, 3–16. New York, NY, USA: ACM.

Bhawalkar, K., and Roughgarden, T. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*.

Bikhchandani, S. 1999. Auctions of heterogeneous objects. *Games and Economic Behavior* 26(2):193 – 220.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; Kyropoulou, M.; Lucier, B.; Paes Leme, R.; and Tardos, É. 2012. On the efficiency of equilibria in generalized second price auctions. *CoRR* abs/1201.6429.

Chan. Y. H, L. L. C. 2012. On linear and semidefinite programming relaxations for hypergraph matching. *Mathematical Programming* 135:123–148.

Christodoulou, G.; Kovács, A.; and Schapira, M. 2008. Bayesian combinatorial auctions. In *ICALP*.

Conitzer, V.; Sandholm, T.; and Santi, P. 2005. Combinatorial auctions with k-wise dependent valuations. In *In Proc. 20th National Conference on Artificial Intelligence (AAAI-05*, 248–254. AAAI Press.

Dobzinski, S., and Schapira, M. 2006. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *SODA*, 1064–1073.

Dobzinski, S.; Nisan, N.; and Schapira, M. 2010. Approximation algorithms for combinatorial auctions with complement-free bidders. *Mathematics of Operations Research* 35:1–13. Preliminary version in STOC'05.

Feige, U., and Izsak, R. 2013. Welfare maximization and the supermodular degree. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, ITCS '13, 247–256. New York, NY, USA: ACM.

Feige, U., and Vondrák, J. 2006. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *FOCS*, 667–676.

Feige, U. 2006. On maximizing welfare when utility functions are subadditive. In *STOC*.

Feldman, M.; Fu, H.; Gravin, N.; and Lucier, B. 2013. Simultaneous auctions are (almost) efficient. In *STOC*.

Gurobi Optimization, I. 2014. Gurobi optimizer reference manual.

Hassidim, A.; Kaplan, H.; Mansour, Y.; and Nisan, N. 2011. Non-price equilibria in markets of discrete goods. In *EC*.

Hazan, E.; Safra, S.; and Schwartz, O. 2006. On the complexity of approximating k-set packing. *Computational Complexity* 15:20–39.

Kelso, A., and Crawford, V. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50:1483–1504.

Lehmann, B.; Lehmann, D.; and Nisan, N. 2001. Combinatorial auctions with decreasing marginal utilities. In *EC*.

Lehmann, D.; O'Callaghan, L.; and Shoham, Y. 1999. Truth revelation in approximately efficient combinatorial auctions. In *EC*, 96–102.

Lucier, B., and Borodin, A. 2010. Price of anarchy for greedy auctions. In *SODA*.

Markakis, E., and Telelis, O. 2012. Uniform price auctions: Equilibria and efficiency. In *SAGT*.

Markov, A. 1890. On a question by D. I. Mendeleev. *Zap. Imp. Akad. Nauk SPb* 62:1–24.

Nisan, N., and Segal, I. 2006. The communication requirements of efficient allocations and supporting prices. *J. Economic Theory* 129:192–224.

Paes Leme, R.; Syrgkanis, V.; and Tardos, E. 2012. Sequential auctions and externalities. In *SODA*.

Roughgarden, T. 2012. The price of anarchy in games of incomplete information. In *EC*.

Sandholm, T. 1999. An algorithm for optimal winner determination in combinatorial auctions. In *IJCAI*, 542–547.

Syrgkanis, V., and Tardos, E. 2013. Composable and efficient mechanisms. In *STOC*.

Syrgkanis, V. 2012. Bayesian games and the smoothness framework. *CoRR* abs/1203.5155.

Vondrák, J. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC*, 67–74.