

Simple Mechanisms for Agents with Complements

MICHAL FELDMAN, Tel Aviv University
OPHIR FRIEDLER, Tel Aviv University
JAMIE MORGENSTERN, University of Pennsylvania
GUY REINER, Tel Aviv University

We study the efficiency of simple auctions in the presence of complements. Devanur et al. [2015] introduced the single-bid auction, and showed that it has a price of anarchy (PoA) of $O(\log m)$ for complement-free (i.e., subadditive) valuations. Prior to our work, no non-trivial upper bound on the PoA of single bid auctions was known for valuations exhibiting complements. We introduce a hierarchy over valuations, where levels of the hierarchy correspond to the degree of complementarity, and the PoA of the single bid auction degrades gracefully with the level of the hierarchy. This hierarchy is a refinement of the Maximum over Positive Hypergraphs (MPH) hierarchy [Feige et al. 2015], where the degree of complementarity d is captured by the maximum number of neighbors of a node in the positive hypergraph representation. We show that the price of anarchy of the single bid auction for valuations of level d of the hierarchy is $O(d^2 \log(m/d))$, where m is the number of items. We also establish an improved upper bound of $O(d \log m)$ for a subclass where every hyperedge in the positive hypergraph representation is of size at most 2 (but the degree is still d). Finally, we show that randomizing between the single bid auction and the grand bundle auction has a price of anarchy of at most $O(\sqrt{m})$ for general valuations. All of our results are derived via the smoothness framework, thus extend to coarse-correlated equilibria and to Bayes Nash equilibria.

General Terms: Mechanism Design, Algorithms, Performance

Additional Key Words and Phrases: Price of Anarchy, Restricted Complements, Auctions

1. INTRODUCTION

A central focus of algorithmic mechanism design is to decide how to allocate limited resources to strategic agents while taking into account computational limitations. A great deal of work has studied truthful mechanisms, and while many times achieving guarantees that match the algorithmic problem (in which the agents are not strategic but always truth telling), many of the designed mechanisms turned out quite complex algorithmically and complicated to describe.

Practical concerns have led recent study to forgo truthfulness in lieu of simple mechanism formats. Simultaneous item auctions (SIAs), in particular, have constant-factor welfare approximations at equilibrium for subadditive buyers [Feldman et al. 2013a], and have an arguably simple format: each buyer submits a single sealed bid for each item separately, and each item's winner is the highest bidder for that item. Unfortunately, SIAs have a marked *lack* of simplicity in another respect: there is initial evidence that the problem of computing Nash equilibria [Dobzinski et al. 2015], approximate Bayes Nash equilibria, correlated equilibria, or verifying best-responses [Cai and Papadimitriou 2014] are likely intractable.

This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. Copyright © 2016 ACM 978-1-4503-3936-0/16/07 ...\$15.00.

<http://dx.doi.org/10.1145/2940716.2940735>

So, while SIAs have a simple format, the strategic behavior induced by the mechanism is quite complex. A mechanism with a simple format but one that is difficult to play leads one to question the underlying assumption that an equilibrium will be reached, and in turn to question the applicability of the price of anarchy bounds.

Recent work [Devanur et al. 2015] introduced another mechanism whose format was “simple” with a strategy space small enough that no-regret learning algorithms (for computing correlated and coarse correlated equilibria of the mechanism) run in polynomial time. This mechanism was coined the *single bid* mechanism, and was shown to have a Price of Anarchy (PoA) of $O(\log m)$ for subadditive buyers, where m is the number of items. This upper bound on the PoA, while worse than that of SIAs, should apply to the welfare achieved by polynomially bounded agents (unlike those for SIAs).

The format of the single-bid mechanism was generalized by Braverman et al. [2016], who defined the notion of a priori learnable interpolation (ALI) mechanisms. An ALI mechanism has two phases. First, agents report $O(\log m)$ bits of information to the mechanism. The mechanism computes some truthful mechanism as a function of all agents’ reports. Second, the agents interact with this truthful mechanism. Since the second interaction is with a truthful mechanism, agents strategize only over their reports in the first phase. To find reports for the first round which form an equilibrium, one can trivially employ no-regret learning in polynomial time over the possible $\text{poly}(m)$ reports. Thus, these mechanisms are strategically simple. If the truthful mechanism selected at the second phase always has a simple format, then the ALI mechanism will also have a simple format.

Both SIAs and single-bid auctions provide good approximation guarantees for complement-free (i.e., subadditive) bidders. However, valuations with *complementarities* arise naturally in many contexts, such as radio spectrum auctions, auctions for landing and takeoff time slots in airports, auctions for computational resources in the cloud, and more (see [Cramton et al. 2006]).

In this work, we aim to design mechanisms for bidders with complementarities, which simultaneously approximate optimal welfare at equilibrium, have a simple format, and are strategically simple (as defined implicitly by Devanur et al. [2015] and formally by Braverman et al. [2016]). Formally, we wish to find mechanisms that run in polynomial time, whose equilibria have high welfare, and whose equilibria can be found in a computationally efficient manner, when bidders’ valuations are not necessarily subadditive.

Several classes of valuations with restricted complements have been proposed in the literature: (1) positive hypergraphs with rank at most k (PH- k), where the valuation is represented by a weighted hypergraph, the hyperedges have positive weights, and are of size at most k . The valuation for a set of items S is the sum of the weights of the hyperedges contained in S . (2) maximum over PH- k (MPH- k), where the value for a set of items S is the maximum value assigned to S across multiple PH- k valuations. (3) supermodular- d (SM- d), where the following graph is considered: the nodes correspond to goods, and an edge (i, j) indicates complementarity between the goods i and j ¹. The complementarity level d corresponds to the maximum degree of any node in the graph.

While for SIAs, the PoA for MPH- k is bounded by $2k$ [Feige et al. 2015], for single bid auctions, the PoA can be linear in m even for PH-2.

Our main result is that a refinement of the MPH- k hierarchy captures the degradation of the single bid auction in the presence of complements. This refinement is described next. The number of neighbors of a node in the PH representation of a valuation is the number of nodes with which a node shares a hyperedge. Our hierarchy, called *Maximum over Positive Supermodular d* (MPS- d), is a maximum over a col-

¹Goods i and j are said to exhibit complementarity if there exists some set S such that $v(j|S \cup i) > v(j|S)$.

lection of PH valuations in which the maximum number of neighbors of any node is bounded by d . Clearly, this constraint implies that each valuation in the collection is also in PH- d . One can verify that this refinement of PH- d is equivalent to the intersection of PH and SM- d . This hierarchy is complete; the highest level of the hierarchy, MPS- m , captures all monotone functions. We show that the price of anarchy of the single bid auction degrades gracefully with the degree of the MPS- d hierarchy. This is cast in the following theorem.

Theorem: When agents have MPS- d valuations, the single-bid auction has a price of anarchy of at most $\frac{(d+1)}{1-e^{-(d+1)}} \cdot (d+2) \cdot H_{\frac{m}{d+1}}$ ($= O(d^2 \log(m/d))$) w.r.t. coarse correlated equilibria².

We also show that for a subclass of MPS- d valuations, where for every PH representation in the collection every hyperedge is of size at most 2 (i.e., where weights are given only to nodes and edges in the graph), and the number of neighbors of a node is bounded by d (i.e., the intersection of PH-2 and SM- d), the price of anarchy is at most $O(d \log m)$.

We prove this by defining an extension of the constrained homogeneous class [Devanur et al. 2015], called d -constrained homogeneous (d -CH), which allows for complements. Our proof proceeds in two steps. We first show that the d -CH class pointwise approximates MPS- d valuations. We then show that the price of anarchy of the single bid auction with d -CH valuations is bounded by $O(d^2 \log(m/d))$.

The above results imply a good price of anarchy for valuations that lie in low levels of the MPS- d hierarchy. We then address the problem of simple auctions for general valuations. We show that randomizing between the single bid auction and the *grand bundle auction* (where the grand bundle is sold via a first price auction) obtains a price of anarchy of at most $O(\sqrt{m})$ for general valuations. Notably, running the two auctions in parallel (by soliciting independent bids) and choosing the better outcome of the two results in a price of anarchy of $\Omega(m)$. The desired result is obtained by randomizing between the two auction formats.

Theorem: The mechanism that randomizes between the single bid auction and the grand bundle auction achieves a price of anarchy of at most $\frac{4\sqrt{m}}{1-e^{-1}}$ for general valuations.

This bound matches the best poly-time welfare approximation by truthful mechanisms (assuming access to a demand oracle) [Dobzinski 2007; Dobzinski et al. 2006; Lavi and Swamy 2005; Nisan and Segal 2006]. It is also known that SIAs cannot achieve a better price of anarchy bound for general valuations [Hassidim et al. 2011]. Finally, it should be noted that there exists a deterministic mechanism that has a price of anarchy of at most $O(\sqrt{m})$ for general valuations [Lucier and Borodin 2010].

1.1. Related work

There has been a great deal of recent focus on simple mechanism design. These mechanisms achieve simplicity of format while trading off the optimality of the allocation they produce; the efficiency of simple, non-truthful mechanisms is measured using the price of anarchy. The goal of this line of research has been to design simple mechanisms whose price of anarchy is as small as possible in as general a setting as possible.

²For ease of exposition H_x denotes the x -th harmonic number when x is an integer and $H_{\lfloor x \rfloor} + 1$ otherwise.

Sequential first-price item auctions have been shown to yield a constant price of anarchy for unit-demand bidders, with respect to subgame perfect equilibrium³ [Leme et al. 2012] and Bayes-Nash equilibria [Syrngkanis and Tardos 2012]. This efficiency breaks for more general classes of valuations than unit-demand bidders: even with one additive bidder and $n-1$ unit-demand bidder, the pure Nash PoA can be $\Omega(m)$ [Feldman et al. 2013b].

The techniques for upper-bounding the Bayes-Nash PoA were shown to be generally useful: if one bounds a mechanism’s PoA using a *smoothness* argument (introduced for auctions by Syrgkanis and Tardos [2013], which is closely related to the smoothness of a game [Roughgarden 2009]), then PoA guarantees naturally extend to coarse correlated equilibria of the complete information game as well as Bayes-Nash equilibria.

The study of simultaneous item auctions was initiated by Christodoulou et al. [2008], who showed that when buyers’ valuations are submodular and i.i.d., the Bayesian PoA of second-price SIAs is at most 2, and that Pure Nash equilibria can be computed in polynomial time in the full-information setting for submodular buyers.

First-price simultaneous item auctions have been studied by Hassidim et al. [2011]. They showed that pure Nash equilibria (when they exist) are fully efficient, but that mixed equilibria can have PoA of $\Omega(\sqrt{m})$ for general valuations. In addition, they showed that the price of anarchy for both coarse correlated equilibria with complete information and Bayes-Nash equilibria is $O(m)$ for general valuations, $O(\log m)$ for subadditive valuations, and $O(1)$ for XOS valuations.

SIAs were then shown by Feldman et al. [2013a] to have constant PoA at Bayes-Nash equilibria for subadditive buyers, for both first and second price payment rules. This result is tempered somewhat by a string of evidence suggesting that the problem of computing Nash equilibria [Dobzinski et al. 2015] (for subadditive bidders), approximate Bayes-Nash equilibria (even for a mix of unit-demand and additive bidders), correlated equilibria, or verifying best-responses [Cai and Papadimitriou 2014] are likely intractable.

Another simple auction format that does allow for efficient computation of its coarse correlated equilibria (using no-regret learning algorithms and demand oracles) is the single-bid auction. In this auction, each bidder submits a single real number, and buyers (in descending order of their bids) choose a bundle amongst the remaining items, paying their bid for each item. This auction format was introduced by Devanur et al. [2015], where the authors showed its price of anarchy of $O(\log m)$ for coarse correlated equilibria with subadditive bidders. The computational efficiency relied on the mechanism having a single round of strategic play which has a small action space, followed by a round of truthful behavior where agents select a utility-maximizing bundle. Braverman et al. [2016] showed that this was essentially the best welfare one could achieve using any interpolation protocol which first has a single round of strategic play over a small action space, followed by some non-adaptive posted price mechanism.

Lucier and Borodin [2010] give a deterministic greedy mechanism that has $O(\sqrt{m})$ price of anarchy. There are also truthful mechanisms that give $O(\sqrt{m})$ approximation to welfare, and run in poly time given an access to a demand oracle [Dobzinski 2007; Dobzinski et al. 2006; Lavi and Swamy 2005].

Several notions of hierarchical restricted complements have been introduced in the literature. Abraham et al. [2012] introduce positive hypergraph representations of valuations with rank at most k , PH- k , give k -approximation algorithms for welfare approximation and $O(\log^k m)$ -approximate truthful mechanisms for this class (and show the algorithmic result is the best possible in polynomial time unless $P = NP$). Feige

³The natural extension of Nash Equilibrium to sequential games.

and Izsak [2013] introduce the notion of supermodular degree (at most) d , SM- d . When valuations are in SM- d , they show APX-hardness of answering demand queries for SM- d for $d \geq 3$, and construct two $(d + 2)$ -approximation algorithms for welfare maximization. Feige et al. [2015] introduce a complete hierarchy of monotone functions, the maximum over positive hypergraphs with rank at most k , MPH- k . They give a $(k + 1)$ -approximation to welfare maximization for this class, and show that SIAs have a price of anarchy at most $2k$ when buyers' valuations are contained in MPH- k .

Simple auction design has also been studied in the context of revenue maximization, both in single-parameter [Devanur et al. 2011; Dhangwatnotai et al. 2010; Hartline and Roughgarden 2009; Morgenstern and Roughgarden 2015] and multiparameter [Babaioff et al. 2014; Chawla et al. 2007, 2010; Rubinstein and Weinberg 2015; Yao 2015] contexts.

Organization of the paper. In section 2 we formally define the setting and introduce our refined hierarchy, MPS- d . In section 3 we prove our main results, namely upper and lower bounds on the price of anarchy of the single bid auction with MPS- d valuations. In section 4 we prove that randomizing between the single bid auction and the grand bundle auction achieves a price of anarchy of $O(\sqrt{m})$ for general valuations. We conclude in Section 5 with a discussion and some open problems. Due to space limitations, some of the proofs are deferred to the full version [Feldman et al. 2016].

2. PRELIMINARIES

A combinatorial auction design problem consists of a set N of n agents, and a set of goods $[m] = \{1, 2, \dots, m\}$. Each agent i has a private valuation function $v_i : 2^{[m]} \rightarrow \mathbb{R}_+$. We use \mathbf{v} to denote the valuation profile $(v_i)_{i \in N}$. We also write $\mathbf{v} = (v_i, \mathbf{v}_{-i})$, where \mathbf{v}_{-i} denotes the valuations of all agents other than i . We design auctions which allocate each agent i a set of goods S_i , such that the social welfare $\text{SW}(S) = \sum_i v_i(S_i)$ is (approximately) maximized. Let $\text{OPT}(\mathbf{v})$ be an allocation that maximizes the social welfare for the valuation profile \mathbf{v} . Fixing an auction and the behavior of all n agents, each agent is charged some payment $P_i \geq 0$. An agent i with valuation v_i who is allocated a set of items S_i and charged P_i has quasi-linear utility $u_i = v_i(S) - P_i$. We will assume agents will behave to maximize this utility.

A mechanism is *truthful* if truth-telling is a dominant strategy; i.e., each agent maximizes its utility by reporting truthfully, regardless of its valuation and other agents' actions. An *interpolation mechanism* is a communication protocol with two phases. The first phase is non-truthful, and its output is a truthful mechanism.

Definition 2.1. (Braverman et al. [2016]) An interpolation mechanism is *a priori learnable* if the first phase contains a single simultaneous broadcast round of communication, and the per-agent communication is $O(\log m)$.

The following observation describes the key property that motivates the study of a priori learnable interpolation (ALI) mechanisms.

OBSERVATION 2.2. [Braverman et al. 2016] *An agent can run a regret-minimizing algorithm over her strategies in an a priori learnable interpolation mechanism (ALI) in time/space poly(m). Therefore, a correlated equilibrium of any ALI can be found in poly-time, and correlated equilibria arise as the result of poly-time distributed regret minimization.*

The Single-bid auction. The single-bid auction, recently introduced by Devanur et al. [2015], is an ALI mechanism. In the first phase the auctioneer solicits a single bid $b_i \in \mathbb{R}_+$ from each agent i . In the second phase the auctioneer sequentially approaches the agents, in a decreasing order of their bids (ties are broken arbitrarily), and offers

each agent i to purchase any of the items that have not been purchased yet, at a per-item price of b_i . We assume that agents maximize their utility: when offered a set of items $U \subseteq [m]$, agent i selects a set $S_i \in \arg \max_{S \subseteq U} v_i(S_i) - |S_i| \cdot b_i$. Notice that fixing the first phase of the single-bid auction, the second phase is truthful; that is, reporting a set in $\arg \max_{S \subseteq U} \{v_i(S_i) - |S_i| \cdot b_i\}$ maximizes utility. Therefore, we assume that agent i behaves strategically only when reporting her bid in the first phase, and truthfully selects a utility-maximizing set in the second phase. Assuming that a single bid can be expressed using communication size of $O(\log m)$, the single bid auction is an ALI mechanism.

Price of Anarchy and smoothness. The allocation resulting from strategic play in the single-bid auction can result in a sub-optimal allocation of goods. Observation 2.2 implies that agents employing no-regret algorithms will converge to an (approximate) correlated or coarse correlated equilibrium. Therefore, it is of interest to provide efficiency guarantees on correlated and coarse equilibria. This efficiency is measured via the *price of anarchy* (PoA), which is the ratio of the optimal social welfare to the welfare at the worst possible equilibrium. Given an equilibrium eq , denote by $\text{SW}(eq)$ the social welfare at this equilibrium.

Definition 2.3. Let E denote any solution concept for mechanism \mathcal{M} , and let \mathcal{V} be a class of valuation profiles. Then the price of anarchy (PoA) and the price of stability (PoS) of \mathcal{M} with respect to E when the agents' valuation profile is in \mathcal{V} are:

$$PoA = \max_{\mathbf{v} \in \mathcal{V}} \max_{eq \in E} \frac{\text{SW}(OPT(\mathbf{v}))}{\text{SW}(eq)} \quad PoS = \max_{\mathbf{v} \in \mathcal{V}} \min_{eq \in E} \frac{\text{SW}(OPT(\mathbf{v}))}{\text{SW}(eq)}$$

All our positive results apply to coarse correlated equilibria and Bayes-Nash equilibria.

Definition 2.4. (Coarse Correlated Equilibrium) An α -coarse correlated equilibrium is a joint distribution σ over bid vectors, such that for each agent i and bid b'_i :

$$\mathbb{E}_{\mathbf{b} \sim \sigma} [u_i(\mathbf{b})] \geq \mathbb{E}_{\mathbf{b} \sim \sigma} [u_i(b'_i, \mathbf{b}_{-i})] - \alpha$$

The smoothness notion was introduced by Roughgarden [2009] for general games, and Syrgkanis and Tardos [2013] applied it to mechanisms. The smoothness framework provides a method for proving price of anarchy upper bounds for various solution concepts.

Definition 2.5. (Syrgkanis and Tardos [2013]) A mechanism \mathcal{M} is (λ, μ) -smooth for a class of valuations $\mathcal{V} = \times_i \mathcal{V}_i$ if for any valuation profile $\mathbf{v} \in \mathcal{V}$, there exists a (possibly randomized) action profile $\mathbf{a}_i^*(\mathbf{v})$ such that for every action profile \mathbf{a} :

$$\sum_i \mathbb{E}_{\mathbf{a}'_i \sim \mathbf{a}_i^*(\mathbf{v})} [u_i(\mathbf{a}'_i, \mathbf{a}_{-i}; v_i)] \geq \lambda \cdot \text{SW}(OPT(\mathbf{v})) - \mu \sum_i P_i(\mathbf{a}) \quad (1)$$

THEOREM 2.6. (Syrgkanis and Tardos [2013]) *If a mechanism is (λ, μ) -smooth then the price of anarchy w.r.t. coarse correlated equilibria is at most $\frac{\max\{1, \mu\}}{\lambda}$.*

2.1. Categories of valuation functions

A set function $f : 2^{[m]} \rightarrow \mathbb{R}_+$ is *normalized* if $f(\emptyset) = 0$ and *monotone* if $f(T) \leq f(S)$ for every $T \subseteq S$. As standard, we assume that all valuations are normalized and monotone.

A *hypergraph representation* of a set function f is a (normalized, but not necessarily monotone) set function h such that for every set $S \subseteq [m]$ it holds that $f(S) = \sum_{T \subseteq S} h(T)$. One can easily verify that every set function f has a unique hypergraph representation h .

A set function is complement-free, or subadditive, if for all $S, T \subseteq [m]$ it holds that $f(S \cup T) \leq f(S) + f(T)$.

When studying a class of valuations \mathcal{V} , it can be useful to also study the class $\max(\mathcal{V})$, as defined below.

Definition 2.7. Given a class of valuations \mathcal{V} , the class $\max(\mathcal{V})$ is the class of all valuations that can be represented as a maximum over a collection of valuations from \mathcal{V} , i.e., $\max(\mathcal{V}) = \{f : \exists \mathcal{G} \subseteq \mathcal{V} : \forall S \subseteq [m], f(S) = \max_{g \in \mathcal{G}} g(S)\}$.

In this paper we focus on valuation functions that exhibit complements. The following hierarchies of valuations with complements have been considered in the literature.

Maximum over positive hypergraphs. [Feige et al. 2015] The class PH (*positive-hypergraph*) is the class of all functions f whose hypergraph representation h has nonnegative edges. The class PH- k contains all functions $f \in \text{PH}$ for which every set T with $h(T) > 0$ satisfies $|T| \leq k$. The class *maximum over PH- k* (MPH- k) is the class $\max(\text{PH-}k)$. Unlike PH- k , MPH- k is a *complete hierarchy*: for every set function f , there exists some $k \leq m$ such that f is in MPH- k (in particular, all functions are in MPH- m).

The supermodular degree. [Feige and Izsak 2013] The *supermodular degree* measures the extent to which any set function f exhibits supermodular behavior. For an item j and set S , denote by $f(j|S) = f(S \cup j) - f(S)$ ⁴ the marginal value of item j given S . The supermodular dependency set of item j is defined as $\text{Dep}^+(j) = \{j' : \exists S \subseteq [m] \text{ so that } f(j|S \cup j') > f(j|S)\}$. The supermodular degree of f is defined as $\max_{j \in [m]} |\text{Dep}^+(j)|$. The class *supermodular degree d* (SM- d) contains all the set functions with supermodular degree at most d . Clearly, the SM- d hierarchy is complete, as any set function has supermodular degree at most $m - 1$.

2.2. A refined hierarchy of restricted complements

The lowest level in the MPH- k hierarchy (MPH-1) is contained in the class of subadditive valuations. It follows from Devanur et al. [2015] that for MPH-1 valuations the price of anarchy is upper bounded by $\frac{e}{e-1} H_m$ (where H_m is the m 'th harmonic number). However, this positive result does not extend beyond the lowest level of MPH (or even PH).

OBSERVATION 2.8. [Morgenstern 2015] *The single bid auction has price of stability of at least m when agents have valuations in PH-2.*

PROOF. A t -star-graph, centered at j , is a graph with t nodes, where there is an edge between the center node (j) and each one of the other $t - 1$ nodes. A *t -star-shaped valuation* is a valuation with a t -star-graph hypergraph representation, in which all edges have weight 1.

Consider two agents, a and b , and the items $[m]$. Let v_a be an m -star-shaped valuation, centered at item 1. Therefore, for all $T \subseteq [m]$, $v_a(T) = |T| - 1$ if $1 \in T$ and 0 otherwise. By construction, $v_a \in \text{PH-2}$. Agent b only wants item 1 for a value of $(m - 1)/m + \epsilon$. For agent a to purchase item 1 in equilibrium, it must pay at least $(m - 1)/m + \epsilon$, otherwise, agent b can bid slightly higher than a 's bid and improve its utility. However, if agent a acquires a set $T \ni 1$ for a price p per item, its utility is $|T| \cdot (1 - p) - 1$. Therefore, if agent a bids more than $(m - 1)/m$, buying any set of items yields negative utility. As a result, at any equilibrium, agent b gets item 1, agent a has 0 value, and the social welfare is $(m - 1)/m + \epsilon$. In the optimal outcome, agent a gets all

⁴We abuse notation and write $S \cup j$ instead of $S \cup \{j\}$

the items and the social welfare is $m - 1$. Therefore, the fraction of the optimal welfare that is achieved in any pure equilibrium is $\frac{(m-1)/m+\epsilon}{m-1} = 1/m + \frac{\epsilon}{m-1}$. \square

This bound is essentially tight. Indeed, it is easy to show⁵ that the single bid auction is $((1 - e^{-m})/m, 1)$ -smooth for *general valuations*, implying a price of anarchy of at most $m/(1 - e^{-m})$. This example demonstrates that the second level of the MPH hierarchy contains valuations that render the worst possible setting for the single bid auction. While one may interpret this result to imply that single bid auctions are hopeless in the presence of complements, we show that viewing restricted complements using a different lens reveals the effect of the level of complementarity on the performance of the auction. In particular, we establish positive results for a refined hierarchy, which combines the structural properties of both SM- d and MPH- k valuations. One would hope that the SM- d hierarchy, by itself, would enable positive price of anarchy results. This is left as an open problem of this work.

Maximum over Positive-Supermodular- d . Positive Supermodular d (PS- d) functions are functions that have a PH representation, and furthermore their supermodular degree is at most d . Thus, in PS- d valuations each item can have at most d items in its supermodular dependency set (in the sense of [Feige and Izsak 2013]).

Definition 2.9. (Maximum over Positive-Supermodular- d) The class Positive Supermodular (PS- d) is defined as $\text{PS-}d = \text{SM-}d \cap \text{PH}$, and the class MPS- d is defined as $\text{MPS-}d = \max(\text{PS-}d)$.

Lemma 2.10 shows that PS- d can be equivalently defined as the class of valuations with PH representation in which the number of neighbors of every node is at most d . This trivially implies that PS- d is contained in PH- $(d + 1)$.

LEMMA 2.10. *Let v be a valuation in PS- d with a hypergraph representation w . For any two items $j, j' \in [m]$, it holds that $j' \in \text{Dep}^+(j)$ if and only if there exists a hyperedge e for which $w_e > 0$ and $\{j, j'\} \subseteq e$.*

The MPS- d hierarchy is *complete*⁶, i.e., for every monotone valuation f there exists some $d \leq (m - 1)$ such that $f \in \text{MPS-}d$.

3. THE SINGLE BID AUCTION IN THE PRESENCE OF COMPLEMENTARITIES

We now present the main result of this section, namely, that the welfare in any coarse correlated equilibrium of the single-bid auction, when buyers' valuations are in MPS- d , is an approximation to the optimal welfare.

THEOREM 3.1. *For agents with valuations in MPS- d , the coarse correlated price of anarchy of the single-bid auction is no more than $\frac{1}{1 - e^{-(d+1)}}(d + 1)(d + 2) \cdot H_{\frac{m}{d+1}}$.*

Specifically, we show that when agents have MPS- d valuations, the single bid auction is a $(\frac{1 - e^{-(d+1)}}{(d+1) \cdot (d+2) \cdot H_{\frac{m}{d+1}}}, 1)$ -smooth mechanism.

In addition, we prove a stronger upper bound of $\frac{2(d+1)}{1 - e^{-2}} \cdot H_{m/2}$ when agents have $\max(\text{PH-2} \cap \text{SM-}d)$ valuations, which is a strict subclass of MPS- d . We also show a PoS lower bound of $\Omega(d + \frac{\log m}{\log \log m})$ when agents have $\text{PH-2} \cap \text{SM-}d$ valuations.

The following proof method for establishing the smoothness of a mechanism with respect to a class of valuations \mathcal{V} was presented in Devanur et al. [2015]: first show

⁵As a corollary from Lemma 4.5

⁶Since $\text{PS-}(m-1) = \text{SM-}(m-1) \cap \text{PH} = \text{PH}$, we get that $\text{MPS-}(m-1) = \text{MPH-}m$.

smoothness for a restricted class of valuations \mathcal{V}' . Then, show that the class \mathcal{V} can be *pointwise β -approximated* by the restricted class \mathcal{V}' . Pointwise approximation is defined as follows:

Definition 3.2. [Devanur et al. 2015] [*pointwise β -approximation*] A valuation class \mathcal{V} is pointwise β -approximated by a valuation class \mathcal{V}' if for any valuation $v \in \mathcal{V}$ and for any set $S \subseteq [m]$, there exists a valuation $v' \in \mathcal{V}'$ such that $\beta \cdot v'(S) \geq v(S)$ and for all $T \subseteq [m]$ it holds that $v'(T) \leq v(T)$.

Note that pointwise β -approximation is less restrictive than mapping each valuation $v \in \mathcal{V}$ to a single valuation $v' \in \mathcal{V}'$ that approximates it everywhere, yet smoothness of a mechanism for valuations in \mathcal{V}' implies smoothness for the larger class \mathcal{V} .

LEMMA 3.3. [Devanur et al. 2015] *If a mechanism for a combinatorial auction setting is (λ, μ) -smooth for the class of valuations \mathcal{V}' and \mathcal{V} is pointwise β -approximated by \mathcal{V}' , then it is $(\frac{\lambda}{\beta}, \mu)$ -smooth for the class \mathcal{V} .*

A constraint-homogeneous (CH) valuation is an additive valuation such that the value of every item is either 0 or \hat{v} for some fixed $\hat{v} > 0$. In Devanur et al. [2015] it was proved that complement-free valuations are pointwise H_m -approximated by CH valuations.

When trying to apply a similar technique for the case of PS-d valuations, we face a challenge, namely that for $d \geq 1$ PS-d valuations cannot be pointwise β -approximated by complement-free valuations for any β . To see this, consider an instance with two items $\{a, b\}$ and the PS-1 valuation $v(\{a\}) = v(\{b\}) = 0$ and $v(\{a, b\}) = 1$. Any complement-free valuation $v' \leq v$ will have $v'(\{a\}) = v'(\{b\}) = 0$ which implies $v'(\{a, b\}) = 0$. Therefore, in order to use the technique of pointwise approximation for PS-d valuations one must go beyond complement-free valuations. To this end we introduce the following class of valuations.

Definition 3.4. (*d-Constraint Homogeneous Valuations*) A valuation v is *d-constraint homogeneous (d-CH)* if there exists a value \hat{v} , and disjoint sets of items Q_1, \dots, Q_ℓ , each of size at most d , so that $v(Q_i) = \hat{v} \cdot |Q_i|$ for every Q_i , and the value of every set $S \subseteq [m]$ is the sum of values of contained Q_i 's, i.e.,

$$v(S) = \sum_{Q_i \subseteq S} v(Q_i) = \hat{v} \sum_{Q_i \subseteq S} |Q_i| = \hat{v} \cdot |\{t : \exists i \text{ s.t. } t \in Q_i \subseteq S\}|$$

Note that 1-CH valuations are CH valuations and that d -CH valuations contain single minded bidders where the interest set of each agent is of size at most d . The remainder of this section is structured as follows. In Lemma 3.6 we show that when agents have d -CH valuations the single bid auction is a $(\frac{1-e^{-d}}{d}, 1)$ -smooth mechanism. In Lemma 3.7 we show that the class of PS-d valuations is pointwise $(d+2) \cdot H_{\frac{m}{d+1}}$ -approximated by $(d+1)$ -CH valuations. These two lemmas imply the smoothness result for PS-d. Finally, Observation 3.5⁷ implies that the same smoothness result carries over to MPS-d.

OBSERVATION 3.5. *For every valuation class \mathcal{V} , the valuation class $\max(\mathcal{V})$ is pointwise 1-approximated by \mathcal{V} .*

We begin by proving smoothness for agents with d -CH valuations.

⁷Observation 3.5 appeared previously (e.g. Lucier and Syrgkanis [2015]; Syrgkanis and Tardos [2013]) and its proof is by definition: for a valuation $v \in \max(\mathcal{V})$ and a set $S \subseteq [m]$, let $v^* = v_\ell$ so that $\ell \in \arg \max_{\ell \in \mathcal{C}} v_\ell(S)$, then by definition $v(S) = v^*(S)$ and $v(T) \geq v^*(T)$ For any set $T \subseteq [m]$.

LEMMA 3.6. *The single bid auction is a $((1 - e^{-d})/d, 1)$ -smooth mechanism when agents have d -CH valuations.*

PROOF. Fix a valuation profile \mathbf{v} of d -CH valuations, and let $S^* = \text{OPT}(\mathbf{v})$ be an optimal allocation w.r.t. \mathbf{v} . Fix an agent i , let \hat{v} and $\{Q_\ell\}_\ell$ be the parameters in agent i 's valuation, and for presentation clarity write $v = v_i$; $v(S) = \hat{v} \cdot \sum_{Q_\ell \subseteq S} |Q_\ell|$. Consider a bid profile \mathbf{b} and denote by $p_j(\mathbf{b})$ the induced price for item j , i.e., $p_j(\mathbf{b}) = b_{i^*}$ so that i^* is the agent that purchases j under bid profile \mathbf{b} . Consider an arbitrary set $Q_\ell \subseteq S_i^*$. Agent i can acquire all items in Q_ℓ by bidding $t > \max_{j \in Q_\ell} p_j(\mathbf{b})$. In such a case the utility from purchasing Q_ℓ is $v(Q_\ell) - t \cdot |Q_\ell| = \hat{v} \cdot |Q_\ell| - t \cdot |Q_\ell| = |Q_\ell| \cdot (\hat{v} - t)$. Therefore:

$$u_i(t, \mathbf{b}_{-i}) \geq \sum_{Q_\ell \subseteq S_i^*} |Q_\ell| \cdot (\hat{v} - t) \cdot \mathbb{1}\{t > \max_{j \in Q_\ell} p_j(\mathbf{b})\}$$

Suppose i performs the randomized deviation $a_i^*(v_i)$ with the density function $f(t) = \frac{1}{d} \cdot \frac{1}{\hat{v}-t}$ and support $[0, (1 - e^{-d}) \cdot \hat{v}]$, Then:

$$\begin{aligned} \mathbb{E}_{t \sim a_i^*(v_i)} [u_i(t, \mathbf{b}_{-i})] &\geq \sum_{Q_\ell \subseteq S_i^*} |Q_\ell| \cdot \int_{\max_{j \in Q_\ell} \{p_j(\mathbf{b})\}}^{(1-e^{-d})\hat{v}} (\hat{v} - t) \cdot f(t) dt \\ &= \frac{1}{d} \cdot \sum_{Q_\ell \subseteq S_i^*} |Q_\ell| \cdot \left((1 - e^{-d})\hat{v} - \max_{j \in Q_\ell} \{p_j(\mathbf{b})\} \right) \end{aligned}$$

By $\max_{j \in Q_\ell} \{p_j(\mathbf{b})\} \leq \sum_{j \in Q_\ell} p_j(\mathbf{b})$ and $v(Q_\ell) = \hat{v} \cdot |Q_\ell|$ and $|Q_\ell| \leq d$ we get that:

$$\mathbb{E}_{t \sim a_i^*(v_i)} [u_i(t, \mathbf{b}_{-i})] \geq \frac{1-e^{-d}}{d} \cdot \sum_{Q_\ell \subseteq S_i^*} v(Q_\ell) - \sum_{Q_\ell \subseteq S_i^*} \sum_{j \in Q_\ell} p_j(\mathbf{b})$$

Finally, the first sum is exactly agent i 's valuations for S_i^* , and the second sum is at most $\sum_{j \in S_i^*} p_j(\mathbf{b})$ since $\{Q_\ell\}_\ell$ is a partition, therefore:

$$\mathbb{E}_{t \sim a_i^*(v_i)} [u_i(t, \mathbf{b}_{-i})] \geq \frac{1-e^{-d}}{d} \cdot v(S_i^*) - \sum_{j \in S_i^*} p_j(\mathbf{b})$$

Summing over all agents establishes the smoothness property. \square

Note that the class of single-minded bidders with interest sets of size at most d is a special case of d -CH valuations, so Lemma 3.6 implies a corresponding bound on the PoA of SBA with respect to single-minded valuations as well.

Next we show that the class PS- d can be pointwise $(d+2) \cdot H_{\frac{m}{d+1}}$ -approximated by $(d+1)$ -CH valuations⁸. In the proof, we use the following two properties of PS- d valuations: First, two items are in the super-dependency set of each other if and only if they share a hyperedge with a positive weight. Second, the size of the super-dependency set of an item is bounded by the level of the hierarchy. We note that neither the class SM- d nor the class PH- k (for $k \geq 2$) exhibit both properties.

LEMMA 3.7. *The PS- d valuation class is pointwise $(d+2) \cdot H_{\frac{m}{d+1}}$ -approximated by the $(d+1)$ -CH valuation class.*

PROOF. Consider a valuation $v \in \text{PS-}d$, a set $X \subseteq [m]$ and some β to be determined later. Let w be the hypergraph representation of v , i.e., $v(S) = \sum_{T \subseteq S} w_T$. Consider

⁸Our proof method is in the spirit of the proof that subadditive valuations are pointwise H_m -approximated by CH valuations, as appears in Devanur et al. [2015]

ALGORITHM 1: Algorithm 1: Partitioning of set X .

Input: A set $X \subseteq [m]$, access to a valuation function v .**Output:** A partition $\mathcal{Q} = \{Q_\ell\}_\ell$ of X

- 1 $S \leftarrow X$.
 - 2 **for** each ℓ from 1 to $\lceil \frac{m}{d+1} \rceil$ **do**
 - 3 Select a set Q_ℓ in $\arg \max_{\substack{A \subseteq S \\ |A|=d+1}} \{v(A)\}$, or $Q_\ell := S$ if $|S| < d+1$.
 - 4 $S \leftarrow S \setminus Q_\ell$. If $S = \emptyset$ then terminate.
 - 5 **end**
-

the following greedy construction of a partition $\mathcal{Q} = \{Q_\ell\}_\ell$ of the set X : While there are more than $d+1$ items, select a subset of yet unselected $d+1$ items from X , with maximum value (with respect to v). The remaining items form the last subset of the partition. The formal description of the greedy process is given in Algorithm 1.

Let $h_{\mathcal{Q}}$ be the function:

$$h_{\mathcal{Q}}(T) = \frac{v(X)}{|\bigcup_{Q_\ell \subseteq T} Q_\ell|^\beta} \cdot \sum_{Q_\ell \subseteq T} |Q_\ell|$$

Note that for any family of disjoint subsets \mathcal{Q}' each of size at most $d+1$, $h_{\mathcal{Q}'}$ is a $(d+1)$ -CH valuation. It suffices to find some $\mathcal{Q}' \subseteq \mathcal{Q}$ so that $\beta \cdot h_{\mathcal{Q}'}(X) \geq v(X)$ and also $h_{\mathcal{Q}'}(T) \leq v(T)$ for all $T \subseteq [m]$. We will examine a sequence of such functions $h_{\mathcal{Q}'}$, so that if none of them pointwise β -approximates v at X , then this implies an upper bound on β .

Initially consider $S_1 = X$. Since \mathcal{Q} is a partition of S_1 we have that $h_{\mathcal{Q}}(X) = \frac{v(X)}{|X|^\beta} \cdot \sum_\ell |Q_\ell| = \frac{v(X)}{\beta}$, so the first requirement of pointwise β -approximation holds. If $h_{\mathcal{Q}}(T) \leq v(T)$ for all $T \subseteq [m]$ then $h_{\mathcal{Q}}$ pointwise approximates v at $|X|$. Otherwise, there exists some T_1 so that $h_{\mathcal{Q}}(T_1) > v(T_1)$. Since v is monotone $v(\bigcup_{Q_\ell \subseteq T_1} Q_\ell) \leq v(T_1) < h_{\mathcal{Q}}(T_1) = h_{\mathcal{Q}}(\bigcup_{Q_\ell \subseteq T_1} Q_\ell)$ therefore we may assume w.l.o.g. that T_1 is a union of sets from \mathcal{Q} . Iteratively, consider $S_i = S_{i-1} \setminus T_{i-1}$. Since T_{i-1} and S_{i-1} are each a union of sets from \mathcal{Q} , then S_i is also a union of sets from \mathcal{Q} , and $\mathcal{Q}_{S_i} = \{Q_\ell \in \mathcal{Q} : Q_\ell \subseteq S_i\}$ is a partition of S_i . By definition, $h_{\mathcal{Q}_{S_i}}(T) = \frac{v(X)}{|S_i|^\beta} \sum_{Q_\ell \in \mathcal{Q}_{S_i} : Q_\ell \subseteq T} |Q_\ell|$ is a $(d+1)$ -CH valuation, and since \mathcal{Q}_{S_i} is a partition of S_i we get that $h_{\mathcal{Q}_{S_i}}(X) = \frac{v(X)}{\beta}$. If for some i it holds that $h_{\mathcal{Q}_{S_i}}(T) \leq v(T)$ for all $T \subseteq [m]$, then $h_{\mathcal{Q}_{S_i}}$ pointwise β -approximates v at X . Otherwise, at some point the iterative process terminates and we are left with two partitions of the set X : $\{Q_\ell\}_\ell$ and $\{T_i\}_i$, so that every Q_ℓ is a subset of some T_j . Therefore:

$$\sum_\ell v(Q_\ell) \leq \sum_i v(T_i) < \sum_i h_{\mathcal{Q}_{S_i}}(T_i) = \frac{v(X)}{\beta} \sum_i \frac{|T_i|}{|S_i|} \quad (2)$$

where the first inequality is because v has a positive-hypergraph representation, the second inequality is by construction, and the last equality is because every S_i and T_i are unions of subsets from \mathcal{Q} . Denote by $\mathcal{C}(\mathcal{Q})$ the collection of all hyperedges $e \subseteq X$ with $w_e > 0$ so that $e \not\subseteq Q_\ell$ for all ℓ . By construction it holds that $v(X) = \sum_\ell v(Q_\ell) + \sum_{e \in \mathcal{C}(\mathcal{Q})} w_e$. The first sum in the last expression is the total weight of all (hyper)edges that are in the interior of some partition element Q_ℓ . The second is the total weight of all edges that connect at least two partition elements. We establish the following lemma:

LEMMA 3.8. $\sum_{e \in \mathcal{C}(\mathcal{Q})} w_e \leq (d+1) \sum_\ell v(Q_\ell)$

Before proving Lemma 3.8 we show how it is used to conclude the proof. Note that the proof of Lemma 3.8 relies on the properties of the class PS-d. Lemma 3.8 implies $v(X) \leq (d+2) \sum_{\ell} v(Q_{\ell})$. By equation (2) we get: $v(X) < (d+2) \frac{v(X)}{\beta} \sum \frac{|T_i|}{|S_i|}$ therefore $\beta < (d+2) \sum \frac{|T_i|}{|S_i|}$. For ease of exposition assume $|X|$ is divisible by $(d+1)$, which implies that the cardinality of every Q_{ℓ} , and hence every S_i and every T_i are divisible by $d+1$. Let $s_i = \frac{|S_i|}{d+1}$ and $t_i = \frac{|T_i|}{d+1}$. Therefore:

$$\sum_i \frac{|T_i|}{|S_i|} = \sum_i \frac{t_i}{s_i} = \sum_i \sum_{j=0}^{t_i-1} \frac{1}{s_i} \leq \sum_i \sum_{j=0}^{t_i-1} \frac{1}{s_i-j} = \sum_{j=0}^{s_1-1} \frac{1}{s_1-j} = H_{s_1} = H_{\frac{|X|}{d+1}} \quad (3)$$

Which concludes that $\beta < (d+2) \cdot H_{\frac{m}{d+1}}$. It remains to prove Lemma 3.8.

PROOF OF LEMMA 3.8. For each Q_{ℓ} , we show there exists a set $E_{\ell} \subseteq \mathcal{C}(\mathcal{Q})$, such that the collection $\{E_{\ell}\}_{\ell}$ satisfies $\mathcal{C}(\mathcal{Q}) \subseteq \cup_{\ell} E_{\ell}$, and for every ℓ it holds that:

$$\sum_{e \in E_{\ell}} w_e \leq (d+1)v(Q_{\ell}) \quad (4)$$

We conclude that $\sum_{e \in \mathcal{C}(\mathcal{Q})} w_e \leq \sum_{\ell} \sum_{e \in E_{\ell}} w_e \leq (d+1) \sum_{\ell} v(Q_{\ell})$, where the first inequality is true since $\mathcal{C}(\mathcal{Q}) \subseteq \cup_{\ell} E_{\ell}$. Let E_{ℓ} denote the set of hyperedges $e \in \mathcal{C}(\mathcal{Q})$ such that ℓ is the minimal index of a set from the partition \mathcal{Q} for which $e \cap Q_{\ell} \neq \emptyset$. For every item $j \in Q_{\ell}$ define $E_{\ell}^j = \{e \in E_{\ell} : j \in e\}$, i.e., the hyperedges in E_{ℓ} in which j is a member, clearly $E_{\ell} = \bigcup_{j \in Q_{\ell}} E_{\ell}^j$. For a set of hyperedges E , let $V(E) = \bigcup_{e \in E} e$. By Lemma 2.10 we get that $V(E_{\ell}^j) \subseteq (\text{Dep}^+(j) \cup \{j\})^9$, which implies that $|V(E_{\ell}^j)| \leq |\text{Dep}^+(j)| + 1 \leq (d+1)$, where the last inequality follows from $\text{PS-d} \subseteq \text{SM-d}$. By definition of E_{ℓ} , for every $j' \in V(E_{\ell})$, if $j' \in Q_{\ell'}$, then $\ell' \geq \ell$, which implies that prior to the ℓ^{th} iteration of step 3 in Algorithm 1, all the items in $V(E_{\ell})$ are available, i.e., in the set S . Therefore, for every item $j \in Q_{\ell}$ the set $V(E_{\ell}^j)$ was available. By step 3 and monotonicity of v , Q_{ℓ} maximizes value over all available sets of size at most $d+1$ therefore $v(Q_{\ell}) \geq v(V(E_{\ell}^j))$ for every j . Therefore:

$$\sum_{e \in E_{\ell}} w_e \leq \sum_{j \in Q_{\ell}} \sum_{e \in E_{\ell}^j} w_e \leq \sum_{j \in Q_{\ell}} v(V(E_{\ell}^j)) \leq |Q_{\ell}| v(Q_{\ell}) \leq (d+1)v(Q_{\ell})$$

□

We also show that $\text{PH-2} \cap \text{SM-d}$ valuations are pointwise $(d+1)H_{m/2}$ -approximated by 2-CH valuations, implying that the PoA is at most $\frac{2(d+1)H_{m/2}}{1-e^{-2}}$ when agents have valuations in $\max(\text{PH-2} \cap \text{SM-d})$:

THEOREM 3.9. *The single bid auction is a $(\frac{1-e^{-2}}{2(d+1)H_{m/2}}, 1)$ -smooth mechanism for $\max(\text{PH-2} \cap \text{SM-d})$ valuations. Thus, it has a price of anarchy of at most $O(d \log(m))$.*

Proposition 3.10 shows a lower bound of d , which holds even for the more restricted class $\text{PH-2} \cap \text{SM-d}$, and even with respect to the best equilibrium.

⁹If $j' \in V(E_{\ell}^j)$ then there exists an edge $e \ni j, j'$ so that $w_e > 0$. By Lemma 2.10 either $j' = j$ or $j' \in \text{Dep}^+(j)$.

PROPOSITION 3.10. *There exists an instance with one bidder with a $SM-d \cap PH-2$ valuation and one bidder that is interested in a single item, for which the price of stability of the single-bid auction is $d - \epsilon$ for every $\epsilon > 0$.*

PROOF. Consider an instance as described in the beginning of subsection 2.2, but with $d + 1$ items. By adding $m - d - 1$ items that have no value to any of the agents, the result follows directly. \square

In [Devanur et al. 2015], a lower bound of $\Omega(\frac{\log m}{\log \log m})$ has been shown for the price of stability (PoS) of the single-bid auction with additive valuations. This bound carries over to valuations in $PH-2 \cap SM-d$ for every d (since additive valuations are a strict subclass of $PH-2 \cap SM-d$). We conclude that the PoS for $PH-2 \cap SM-d$ valuations is at least $\max(d, \Omega(\frac{\log m}{\log \log m}))$. In the full version we show another example that simultaneously captures the two lower bounds above, i.e., an instance where agents have $PH-2 \cap SM-d$ valuations, for which the PoS of the single bid auction is $\Omega(d + \frac{\log m}{\log \log m})$.

PROPOSITION 3.11. *If all agents have valuations in $PH-2 \cap SM-d$, the PoS of the single bid auction w.r.t. pure Nash equilibria is at least $\Omega(d + \frac{\log m}{\log \log m})$.*

4. HYBRID SINGLE BID MECHANISMS

In this section we give a bound on the price of anarchy for general valuations by randomizing between the single bid auction and the *grand bundle auction*, described as follows.

The Grand bundle auction. The grand-bundle auction solicits a single bid $b_i \in \mathbb{R}_+$ from each agent i , approaches the agents in decreasing order of their bids, and offers each agent i the grand bundle $[m]$ for the price b_i , once an agent acquires $[m]$ the auction ends. Since the grand bundle auction solicits a single real-valued bid from each bidder, then runs a truthful mechanism, it is also an ALI mechanism.

We show that randomizing between the single bid auction and the grand bundle auction (by soliciting independent bids for the two auctions) gives price of anarchy of at most $O(\sqrt{m})$ for general valuations (Theorem 4.3).

Our proof proceeds as follows. We partition the space of valuation profiles into two disjoint subspaces, depending on whether there exists an optimal allocation where a single agent contributes at least $\frac{1}{\sqrt{m}}$ of OPT, or not. For the first subspace, we show that the grand bundle auction is $(\frac{1}{2\sqrt{m}}(1 - e^{-1}), 1)$ -smooth; for the second subspace we show that the single bid auction is $(\frac{1}{2\sqrt{m}} \cdot (1 - e^{-\sqrt{m}}), 1)$ -smooth. This immediately implies that running both auctions independently and taking the better equilibrium gives price of anarchy of at most $O(\sqrt{m})$. However, a priori, it is not clear which auction to run. A naive approach would be to run both auctions simultaneously with independent bids, and apply the outcome of the better mechanism. As previously observed by [Lucier and Borodin 2010], such an approach fails badly, as demonstrated by the following example:

Example. Consider $m = n \geq 2$ agents. Each agent i has value 1 for a unique item a_i , and a value of $1 + \epsilon$ for the grand bundle $[m]$. If each agent submits a bid of 0 to the single-bid auction and a bid of $1 + \epsilon$ to the grand bundle auction, the social welfare achieved by the grand bundle auction is $1 + \epsilon$ by handing $[m]$ to some agent. This is a pure Nash equilibrium since no bidder can achieve non-negative utility in the grand bundle auction by deviating, and any unilateral deviation in the single bid auction that achieves non-negative utility for an agent, induces an outcome with social welfare at most 1, therefore the grand bundle auction will still be selected. Clearly the

optimal allocation hands each agent i its respective item of interest a_i , resulting in social welfare of m . This shows a price of anarchy of $m - \epsilon$ for every $\epsilon > 0$.

In [Lucier and Borodin 2010] it is shown that the problem can be circumvented by introducing randomization. In this section we formalize this idea using the smoothness framework.

Definition 4.1. (Hybrid mechanism) Given two mechanisms \mathcal{M} and \mathcal{M}' , and a real number $0 < p < 1$, the *hybrid mechanism* $(\mathcal{M}, \mathcal{M}', p)$ solicits from each agent i two actions, a_i, a'_i , and runs $\mathcal{M}(\mathbf{a})$ with probability p and $\mathcal{M}'(\mathbf{a}')$ with probability $1 - p$.

It follows by definition that a hybrid mechanism that is composed of two ALI mechanisms is also an ALI mechanism.

The following lemma establishes the smoothness of a hybrid mechanism.

LEMMA 4.2. *Let \mathcal{V} and \mathcal{V}' be spaces of valuation profiles. Suppose mechanism \mathcal{M} is (λ, μ) -smooth w.r.t. valuation profiles in \mathcal{V} , and mechanism \mathcal{M}' is (λ', μ') -smooth w.r.t. valuation profiles in \mathcal{V}' . Then, for every p , the hybrid mechanism $(\mathcal{M}, \mathcal{M}', p)$ is $(p \cdot \lambda, \max\{\mu, 1\})$ -smooth w.r.t. valuation profiles in \mathcal{V} and $((1 - p) \cdot \lambda', \max\{\mu', 1\})$ -smooth w.r.t. valuation profiles in \mathcal{V}' .*

PROOF. Consider a valuation profile $\mathbf{v} \in \mathcal{V}$. Consider an arbitrary action profile $(\mathbf{a}, \mathbf{a}')$, where $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{a}' = (a'_1, \dots, a'_n)$. Let P_i and P'_i denote the payments of mechanisms \mathcal{M} and \mathcal{M}' respectively, and similarly for utilities and values. Utilities (u_i^p) , values (v_i^p) , and payments (P_i^p) , denote the expected value of those quantities for agent i in the hybrid mechanism $(\mathcal{M}, \mathcal{M}', p)$ (e.g. for payments, $P_i^p(\mathbf{a}, \mathbf{a}') = p \cdot P_i(\mathbf{a}) + (1 - p) \cdot P'_i(\mathbf{a}')$). Let $a_i^*(\mathbf{v})$ be the deviation given by the smoothness of mechanism \mathcal{M} . For ease of exposition denote $a_i^*(\mathbf{v})$ by a_i^* and assume that a_i^* is a pure strategy. We have that:

$$\sum_i u_i^p((a_i^*, \mathbf{a}_{-i}), \mathbf{a}') = \sum_i p \cdot u_i(a_i^*, \mathbf{a}_{-i}) + (1 - p)u'_i(\mathbf{a}') = p \sum_i u_i(a_i^*, \mathbf{a}_{-i}) + (1 - p) \sum_i u'_i(\mathbf{a}')$$

By smoothness of \mathcal{M} it holds that:

$$\begin{aligned} \sum_i u_i^p((a_i^*, \mathbf{a}_{-i}), \mathbf{a}') &\geq p \cdot \left(\lambda \cdot \text{SW}(\text{OPT}(\mathbf{v})) - \mu \cdot \sum_i P_i(\mathbf{a}) \right) + (1 - p) \cdot \sum_i u'_i(\mathbf{a}') \\ &= p \cdot \lambda \cdot \text{SW}(\text{OPT}(\mathbf{v})) - \mu \cdot p \cdot \sum_i P_i(\mathbf{a}) + (1 - p) \cdot \sum_i (v'_i(\mathbf{a}') - P'_i(\mathbf{a}')) \end{aligned}$$

By $v'_i(\mathbf{a}') \geq 0$ we get:

$$\begin{aligned} \sum_i u_i^p((a_i^*, \mathbf{a}_{-i}), \mathbf{a}') &\geq p \cdot \lambda \cdot \text{SW}(\text{OPT}(\mathbf{v})) - \max\{\mu, 1\} \sum_i (p \cdot P_i(\mathbf{a}) + (1 - p) \cdot P'_i(\mathbf{a}')) \\ &= p \cdot \lambda \cdot \text{SW}(\text{OPT}(\mathbf{v})) - \max\{\mu, 1\} \sum_i P_i^p(\mathbf{a}, \mathbf{a}'), \end{aligned}$$

where the last equality follows by the definition of the hybrid mechanism. Symmetrically, for every valuation profile $\mathbf{v}' \in \mathcal{V}'$ the mechanism is $((1 - p) \cdot \lambda', \max\{\mu', 1\})$ -smooth with respect to valuations in \mathcal{V}' . \square

We now establish our main theorem. While each of the grand bundle and single bid auctions has a price of anarchy of $\Omega(m)$ for general valuations, the hybrid mechanism gives price of anarchy of at most $O(\sqrt{m})$.

THEOREM 4.3. *The hybrid mechanism composed of the single-bid and the grand-bundle auctions with $p = 1/2$ is $(\frac{1}{2\sqrt{m}}(1 - e^{-1}), 1)$ -smooth for general valuations.*

We first consider valuation profiles in which the optimal welfare can be approximated by allocating the grand bundle.

LEMMA 4.4. *If for a class of valuation profiles \mathcal{V} , for every $\mathbf{v} \in \mathcal{V}$ there exists an agent i^* so that $v_{i^*}([m]) \geq \beta \cdot \text{SW}(\text{OPT}(\mathbf{v}))$, then the grand-bundle auction is a $(\beta \cdot (1 - e^{-1}), 1)$ -smooth mechanism.*

We now consider valuation profiles in which the optimal welfare can be well-approximated by “small” allocations.

LEMMA 4.5. *If for every valuation profile \mathbf{v} in a class of valuation profiles \mathcal{V} there exists an allocation S^* so that $\text{SW}(S^*) \geq \beta \cdot \text{SW}(\text{OPT}(\mathbf{v}))$ and $|S_i^*| \leq \gamma$ for every agent i , then for every $c > 0$ the single bid auction is $(c \cdot (1 - e^{-1/c}), \beta \cdot c \cdot \gamma)$ -smooth w.r.t. \mathcal{V} .*

The above lemmas lead to the following definition.

Definition 4.6. A valuation profile \mathbf{v} is z -lopsided if there exists an optimal allocation S^* so that at least half of the social welfare is due to agents that were allocated a bundle with at least z goods, i.e., if $\sum_{i \in A} v_i(S_i^*) \geq \frac{1}{2} \text{SW}(S^*)$, where $A \subseteq N$ and for every $i \in A$ it holds that $|S_i^*| \geq z$. We denote by $LOP(z)$ the class of all z -lopsided valuation profiles.

The following lemma is implied by Lemma 4.4.

LEMMA 4.7. *The grand-bundle auction is a $(\frac{z}{2m} \cdot (1 - e^{-1}), 1)$ -smooth mechanism with respect to valuation profiles in $LOP(z)$.*

Similarly, the following lemma is implied by Lemma 4.5.

LEMMA 4.8. *For every $c > 0$, the single-bid auction is a $(\frac{c}{2} \cdot (1 - e^{-1/c}), c \cdot z)$ -smooth mechanism with respect to valuation profiles $\mathbf{v} \notin LOP(z)$.*

To conclude the proof of Theorem 4.3, note that Lemma 4.7 implies that the grand-bundle auction is $(\frac{1}{2\sqrt{m}}(1 - e^{-1}), 1)$ -smooth w.r.t. valuation profiles in $LOP(\sqrt{m})$, and Lemma 4.8, with $c = \frac{1}{\sqrt{m}}$, implies that the single-bid auction is $(\frac{1}{2\sqrt{m}} \cdot (1 - e^{-\sqrt{m}}), 1)$ -smooth w.r.t. valuations not in $LOP(\sqrt{m})$. It follows by Lemma 4.2 that the hybrid mechanism is $(\frac{1}{4\sqrt{m}}(1 - e^{-1}), 1)$ -smooth, as desired.

We note that this guarantee is tight up to a constant. In the full version we show a lower bound of \sqrt{m} for the PoA of the hybrid mechanism in every PNE.

5. DISCUSSION

In this paper we study simple mechanisms for settings that exhibit complementarities. We focus on the single bid auction, which has been shown to have a logarithmic PoA for complement free valuations. We show upper and lower bounds on the PoA when agents have complementarities, captured by the MPS- d hierarchy. We also show that randomizing between the single bid auction and the grand bundle auction gives PoA of $O(\sqrt{m})$ for general valuations. Our results leave a gap between the lower and upper bounds on the PoA of the single bid auction when applied to MPS- d valuations. In the full version we show that our upper bound is essentially tight with respect to our proof technique. In particular, we show that the pointwise approximation of MPS- d by $(d + 1)$ -CH valuations is tight (up to a $\log m$ factor). A major open problem is to find the price of anarchy of the single bid auction for SM- d valuations.

6. ACKNOWLEDGMENTS

We thank Michael Krivelevich and Asaf Shapira for fruitful discussions.

REFERENCES

- Ittai Abraham, Moshe Babaioff, Shaddin Dughmi, and Tim Roughgarden. 2012. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, New York, NY, USA, 3–16.
- Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. 2014. A simple and approximately optimal mechanism for an additive buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE Computer Society, Washington, DC, USA, 21–30.
- Mark Braverman, Jieming Mao, and S. Matthew Weinberg. 2016. Interpolating Between Truthful and Non-Truthful Mechanisms for Combinatorial Auctions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, Arlington, Virginia, 1444–1457.
- Yang Cai and Christos Papadimitriou. 2014. Simultaneous bayesian auctions and computational complexity. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, New York, NY, USA, 895–910.
- Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 243–251.
- Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, New York, NY, USA, 311–320.
- George Christodoulou, Annamária Kovács, and Michael Schapira. 2008. Bayesian combinatorial auctions. In *Automata, Languages and Programming*. Springer-Verlag, Berlin, Heidelberg, 820–832.
- Peter C Cramton, Yoav Shoham, and Richard Steinberg. 2006. *Combinatorial auctions*. Vol. 475. MIT press Cambridge, Cambridge, MA, USA.
- Nikhil Devanur, Jason Hartline, Anna Karlin, and Thach Nguyen. 2011. Prior-independent multi-parameter mechanism design. In *Internet and Network Economics*. Springer, Berlin, Heidelberg, 122–133.
- Nikhil Devanur, Jamie Morgenstern, Vasilis Syrgkanis, and S Matthew Weinberg. 2015. Simple auctions with simple strategies. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, New York, NY, USA, 305–322.
- Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. 2010. Revenue maximization with a single sample. In *Proceedings of the 11th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 129–138.
- Shahar Dobzinski. 2007. Two Randomized Mechanisms for Combinatorial Auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer-Verlag, Berlin, Heidelberg, 89–103.
- Shahar Dobzinski, Hu Fu, and Robert Kleinberg. 2015. On the complexity of computing an equilibrium in combinatorial auctions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, New York, NY, USA, 110–122.
- Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2006. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. ACM, New York, NY, USA, 644–652.
- Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. 2015. A Unifying Hierarchy of Valuations with Complements and Substitutes. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*. AAAI Press, San Jose, CA, 872–878.
- Uriel Feige and Rani Izsak. 2013. Welfare maximization and the supermodular degree. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. ACM, New York, NY, USA, 247–256.
- Michal Feldman, Ophir Friedler, Jamie Morgenstern, and Guy Reiner. 2016. Simple Auctions For Agents With Complements. *arXiv preprint arXiv:1603.07939* (2016).

- Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. 2013a. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, New York, NY, USA, 201–210.
- Michal Feldman, Brendan Lucier, and Vasilis Syrgkanis. 2013b. Limits of efficiency in sequential auctions. In *Web and Internet Economics*. Springer, Berlin, Heidelberg, 160–173.
- Jason D Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 225–234.
- Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. 2011. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 295–296.
- Ron Lavi and Chaitanya Swamy. 2005. Truthful and near-optimal mechanism design via linear programming. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*. IEEE, Washington, DC, USA, 595–604.
- Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. 2012. Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, New York, NY, USA, 869–886.
- Brendan Lucier and Allan Borodin. 2010. Price of Anarchy for Greedy Auctions. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*. ACM, New York, NY, USA, 537–553.
- Brendan Lucier and Vasilis Syrgkanis. 2015. Greedy algorithms make efficient mechanisms. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, New York, NY, USA, 221–238.
- Jamie Morgenstern. 2015. *Market Algorithms: Incentives, Learning and Privacy*. Ph.D. Dissertation. Carnegie Mellon University.
- Jamie H Morgenstern and Tim Roughgarden. 2015. On the Pseudo-Dimension of Nearly Optimal Auctions. In *Advances in Neural Information Processing Systems*. MIT Press, Cambridge, MA, USA, 136–144.
- Noam Nisan and Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129, 1 (2006), 192–224.
- Tim Roughgarden. 2009. Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*. ACM, New York, NY, USA, 513–522.
- Aviad Rubinfeld and S Matthew Weinberg. 2015. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, New York, NY, USA, 377–394.
- Vasilis Syrgkanis and Eva Tardos. 2012. Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, New York, NY, USA, 929–944.
- Vasilis Syrgkanis and Eva Tardos. 2013. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, New York, NY, USA, 211–220.
- Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, New York, NY, USA, 92–109.