

Combinatorial Auctions without Money*

Dimitris Fotakis
National Technical Univ. of Athens
Greece
fotakis@cs.ntua.gr

Piotr Krysta
University of Liverpool, UK
pkrysta@liverpool.ac.uk

Carmine Ventre
Teesside University, UK
c.ventre@tees.ac.uk

ABSTRACT

Algorithmic Mechanism Design attempts to marry computation and incentives, mainly by leveraging monetary transfers between designer and selfish agents involved. This is principally because in absence of money, very little can be done to enforce truthfulness. However, in certain applications, money is unavailable, morally unacceptable or might simply be at odds with the objective of the mechanism. For example, in Combinatorial Auctions (CAs), the paradigmatic problem of the area, we aim at solutions of maximum social welfare, but still charge the society to ensure truthfulness. We focus on the design of incentive-compatible CAs without money in the general setting of k -minded bidders. We trade monetary transfers with the observation that the mechanism can detect certain lies of the bidders: i.e., we study truthful CAs with verification and without money. In this setting, we characterize the class of truthful mechanisms and give a host of upper and lower bounds on the approximation ratio obtained by either deterministic or randomized truthful mechanisms. Our results provide an almost complete picture of truthfully approximating CAs in this general setting with multi-dimensional bidders.

Categories and Subject Descriptors

F.2.0 [Analysis of Algorithms and Problem Complexity]: General; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Theory, Economics

Keywords

Algorithmic Mechanism Design; Mechanisms with Verification; Combinatorial Auctions

*This work is supported by grant EP/K01000X/1 from the Engineering and Physical Sciences Research Council. D. Fotakis is also supported by AlgoNow project, co-financed by the European Union (European Social Fund) and Greek national funds, through the Operational Program “Education and Lifelong Learning”, under the program THALES.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1. INTRODUCTION

Algorithmic Mechanism Design (AMD) has as main scope the realignment of the objective of the designer with the selfish interests of the agents involved in the computation. One of the principal obstacles to concrete applications of truthful mechanisms is the assumption that mechanisms use monetary transfers. On one hand, money may provoke (unreasonably) large payments [7]; on the other hand, there are contexts for which little justification can be found for either the presence of a digital currency or the use of money at all. There are contexts in which money might be morally unacceptable (such as, to support certain political decisions) or even illegal (as for example, in organ donations). Additionally, there are applications in which the objective of the computation collides with the presence of money.

Consider Combinatorial Auctions (CAs), the paradigmatic problem in AMD. In a CA, we have a set U of m goods and n bidders. Each bidder i has a *private* valuation function v_i that maps subsets of goods to nonnegative real numbers. Agents’ valuations are monotone, i.e., for $S \supseteq T$ we have $v_i(S) \geq v_i(T)$. The goal is to find a partition S_1, \dots, S_n of U such that the *social welfare* $\sum_{i=1}^n v_i(S_i)$ is maximized. For CAs, we are in a paradoxical situation: whilst, on one hand, we pursue the noble goal of maximizing the happiness of the society (i.e., the bidders), on the other, we consider it acceptable to charge the society itself (and then “reduce” its total happiness) to ensure truthfulness. CAs without money would avoid this paradox, automatically guarantee budget-balancedness, and deal with budgeted bidders.

We focus here on k -minded (a.k.a. XOR) bidders, i.e., bidders are interested in obtaining one out of a collection of k subsets of U , and study the feasibility of designing truthful CAs without money, returning reasonable approximations of the optimal social welfare. This is, however, an impossible task in general: it is indeed pretty easy to show that there is no better than n -approximate mechanisms without money, even in the case of single-item auctions and truthful-in-expectation mechanisms [6]. We thus focus on the model of CAs with verification, introduced in [16]. In this model, which is motivated by a number of real-life applications and has also been considered by economists [4], bidders do not overbid their valuations on the set that they are awarded, if any. The hope is that money can be traded with the verification assumption so to be able to design “good” (possibly, polynomial-time) mechanisms, which are truthful without money in a well-motivated, still challenging, model.

Our Contribution. CAs with verification are perhaps best illustrated by means of the following motivating scenario,

discussed first in [16]. Consider a government (auctioneer) auctioning business licenses for a set U of cities under its administration. A company (bidder) wants to get a license for some subset of cities (subset of U) to sell her product stock to the market. Consider the bidder’s profit for a subset of cities S to be equal to a unitary *publicly known* product price (e.g., for some products, such as drugs, the government could fix a social price) times the number of product items available in the stocks that the company possesses in the cities comprising S .¹ In this scenario, the bidder is strategic on her stock availability. As noted, e.g., in [4], a simple inspection on the stock consistency implies that bidders cannot overbid their profits: the concealment of existing items in stock is costless but disclosure of unavailable ones is prohibitively costly. The assumption is therefore a kind of verification *a posteriori*²: the inspection is only carried on for the solution actually implemented, and then each bidder cannot overstate her valuation for the set she gets allocated, if any. It is important to notice that bidders can misreport sets and valuations for unassigned sets in an *unrestricted* way. In similar scenarios, CAs with verification might be run by governments for moneyless bidders (e.g., charities) when the objective is social welfare maximization and an upper bound on the bidders winning valuations can either be obtained by direct inspection or deduced by noticing their features (e.g., their financial situation).

In this model, we first give a complete characterization of truthful mechanisms in both the cases where the collections of k sets, each bidder is interested in, are public (a.k.a. *known bidders*) and private (a.k.a. *unknown bidders*); valuations are always private. The latter distinction is standard in CAs literature (see e.g., [18, 19]), is purely technical and allows the study of truthfulness under different degrees of difficulty. Our results hold either for both settings (i.e., characterizations) or in the harder of the two: unknown bidders for the upper bounds and known bidders for the lower bounds. Truthfulness is characterized in this context in terms of *k-monotone algorithms*: if a known bidder is awarded a set S and augments her declaration for S , then she must get a set “not worse” than S . This generalizes neatly to the case of unknown bidders. Two important considerations follow: (i) Our characterization is significant especially when one observes that the corresponding problem for CAs with money and k -minded bidders is open already for $k = 2$; (ii) Our notions of monotone algorithms generalize the properties of monotonicity shown to characterize truthfulness with money for single-minded bidders [19, 18] and proved to be sufficient for so-called generalized single-minded bidders [2]. This is an interesting development as it is the first case in which a truthful mechanism with money can be “translated” into a truthful mechanism without money. The price to pay is “only” to perform verification to prevent certain lies of the bidders, while algorithms (and their approximation guarantees) remain unchanged. Thus, in light of our results, previously known algorithms presented in, e.g., [18, 2, 10], are truthful not only when money can be used, but also in ab-

¹Bidders will sell products already in stock (i.e., no production costs are involved).

²A stronger model of verification would require bidders to be unable to overbid at all and not just on the awarded set. However, there appears to be weaker motivations for this model: the investment required on inspections would be considerable and rather unrealistic.

sence of money when verification can be implemented. This equivalence gives also a strong motivation for our model.

Armed with the characterization of truthfulness, we provide a number of upper and lower bounds on the approximation guarantee of truthful CAs without money and with verification to the optimal social welfare. The upper bounds hold for the harder case of unknown bidders. In the case where each good in U has a supply b , we give an upper bound of $O(b\sqrt[m]{m})$. This algorithm is deterministic, runs in polynomial time, and adapts an idea of multiplicative update of good prices by [17]. With similar ideas, we obtain randomized universally truthful mechanisms with approximation ratios of $O(d^{1/b} \log(bm))$ and $O(m^{1/(b+1)} \log(bm))$, where d is the maximum size of sets in the bidders’ collections. Our most significant deterministic polynomial-time upper bound is obtained, in the case of $b = 1$, by a simple greedy mechanism that exploits the characteristics of the model without money. This algorithm returns a $\min\{m, d+1\}$ -approximate solution. Two simple randomized universally truthful CAs without money complete the picture: the first achieves a k -approximation in exponential time; the second runs instead in polynomial-time and has a $O(\sqrt{m})$ -approximation guarantee. We note here that all our polynomial-time upper bounds are computationally (nearly) best possible even when the algorithm has full knowledge of the bidders’ data.

We complement this study by showing a host of lower bounds on the approximation guarantee of truthful CAs without money for known bidders, without any computational assumption. We prove the following lower bounds: 2 for deterministic mechanisms; 5/4 for universally truthful mechanisms; and, finally, 1.09 for truthful-in-expectation mechanisms. This implies that the optimal mechanisms are not truthful in our model. Stronger lower bounds are proved for deterministic truthful mechanisms that use priority algorithms [1]. These algorithms process (and take decisions) one *elementary item* at the time, from a list of ordered items. The ordering can also change adaptively after each item is considered (e.g., our greedy mechanism falls in the category of non-adaptive priority algorithms). We give a lower bound of d for priority algorithms that process bids as elementary items (thus, essentially matching the upper bound of the greedy algorithm), and a lower bound of $m/2$ in the case in which the algorithm processes bidders as items.

Our bounds give a very interesting picture of the relative power of verification versus money. For example, we have a $O(\sqrt{m})$ -approximate universally truthful mechanism, which matches the guarantee of the universally truthful mechanism with money given by [5]. However, our lower bounds show that it is not possible to implement the optimal outcome without money; while we can do so in exponential time using VCG payments. If we restrict to polynomial-time mechanisms, we have a deterministic greedy mechanism that is truthful without money and $\min\{m, d+1\}$ -approximate; with money, instead, it is not known how to obtain any polynomial-time deterministic truthful mechanism with an approximation ratio better than $O(m/\sqrt{\log m})$ [13]. Moreover, [1] proved a lower bound of $\Omega(m)$ on the approximation ratio of any truthful greedy mechanism with money for instances with $d = 2$. Our greedy mechanism achieves an approximation ratio of 3 for such instances, which implies that the lower bound of [1] does not hold in our model without money. Additionally, we show that the greedy mechanism cannot be made truthful with money, which suggests that

the model without money couples better with greedy selection rules. A general lower bound in terms of m for CAs without money would shed further light on this debate of power of verification versus power of money; we offer an interesting conjecture in § 5.2. Due to space constraints, we defer some of the proofs to the full version of this paper [8].

Related Work. CAs as an optimization problem (without strategic consideration) is **NP**-hard to solve optimally or even to approximate: neither an approximation ratio of $m^{1/2-\epsilon}$, for any constant $\epsilon > 0$, nor of $O(d/\log d)$ can be obtained in polynomial time [20, 18, 12]. Hence, a large body of literature has focused on the design of polynomial-time truthful CAs that return as good an approximate solution as possible, under assumptions (i.e., restrictions) on bidders' valuation domains, such as single-minded domains [18, 19, 2] and the settings listed in [21, Fig. 11.2].

In [16], instead of restricting the domains of the bidders, it is proposed to restrict the way bidders lie. We adopt here their model, adapting it to the case without money. The definition of CAs with verification is inspired by the literature on mechanisms with verification (see, e.g., [22, 23] and references therein) and by similar models studied in theoretical economics, see, e.g., [11, 4, 3]. The economic model closest to ours is that of [4], where verification takes place for every outcome, and not just for the implemented solution, and is therefore stronger and less realistic than ours. Moreover, the results in [3] suggest that “one-sided” verification in necessary, for otherwise truthful implementation is equivalent to truthful implementation with verification.

Our work fits in the framework of approximate mechanism design without money, initiated by [24]. The idea is that for optimization problems where the optimal solution cannot be truthfully implemented without money, one may resort to the notion of approximation, and seek for the best approximation ratio achievable by truthful mechanisms. Approximate mechanisms without money have been obtained for various problems, among them, for locating 1 or 2 facilities in metric spaces (see e.g., [24]). Due to the apparent difficulty of truthfully locating 3 or more facilities with a good approximation guarantee, notions conceptually similar to our notion of verification have been proposed [9]. Truthful mechanisms without money for scheduling selfish machines whose execution times can be (strongly) verified are considered in [14]. Mechanisms without money for a so-called Generalized Assignment problem are studied in [6]: n selfish jobs compete to be processed by m unrelated machines; the only private data of each job is the set of machines by which it can be actually processed. This problem can be cast as a special case of CAs with $d = 1$ and then [6, Algorithm 1] can be regarded as a special case of our greedy algorithm.

2. MODEL AND PRELIMINARIES

In CAs we have a set U of m goods and n agents, a.k.a. bidders. Each k -minded (a.k.a. XOR-bidder) i has a *private* valuation function v_i and is interested in obtaining only one set in a *private* collection \mathcal{S}_i of subsets of U , and $|\mathcal{S}_i| = k$.³ Each valuation function v_i maps subsets of goods to non-

negative real numbers ($v_i(\emptyset)$ is normalized to be 0), and is monotone: for $S \supseteq T$ we have $v_i(S) \geq v_i(T)$. The goal is to find a partition S_1, \dots, S_n of U such that $\sum_{i=1}^n v_i(S_i)$, the *social welfare*, is maximized.

E.g., consider $U = \{1, 2, 3\}$ and the first bidder to be interested in $\mathcal{S}_1 = \{\{1\}, \{2\}, \{1, 2\}\}$. The valuation function of bidder i for $S \notin \mathcal{S}_i$ is $v_i(S) = \max_{S' \in \mathcal{S}_i: S \supseteq S'} \{v_i(S')\}$ if $\exists S' \in \mathcal{S}_i \wedge S \supseteq S'$, and 0 otherwise. We say that $v_i(S) \neq 0$ (for $S \notin \mathcal{S}_i$) is *defined* by an inclusion-maximal set $S' \in \mathcal{S}_i$ such that $S' \subseteq S$ and $v_i(S') = v_i(S)$. If $v_i(S) = 0$, we say that \emptyset defines it. So, here, $v_1(\{1, 2, 3\})$ is defined by $\{1, 2\}$.

We assume that bidders are interested in sets of size at most $d \in \mathbb{N}$, i.e., $d = \max\{|S| : \exists i \text{ s.t. } S \in \mathcal{S}_i \wedge v_i(S) > 0\}$. We let \mathcal{T}_i be a set of k non-empty subsets of U and z_i be the corresponding valuation of bidder i , i.e., $z_i : \mathcal{T}_i \rightarrow \mathbb{R}_0^+$. We call $b_i = (z_i, \mathcal{T}_i)$ a *declaration* (or *bid*) of bidder i . We let $t_i = (v_i, \mathcal{S}_i)$ be the *true type* of bidder i and let D_i denote the *domain* of bidder i , i.e., the set of all i 's possible bids for all possible types of i .

Fix the bids \mathbf{b}_{-i} of all agents but i . For any $b_i = (z_i, \mathcal{T}_i)$ in D_i , let $A_i(b_i, \mathbf{b}_{-i})$ be the set that auction A on input $\mathbf{b} = (b_i, \mathbf{b}_{-i})$ allocates to bidder i . If no set is allocated to i , we set $A_i(b_i, \mathbf{b}_{-i}) = \emptyset$. We say that A is a truthful auction without money if for any bidder i , $b_i \in D_i$ and \mathbf{b}_{-i} we have:

$$v_i(A_i(t_i, \mathbf{b}_{-i})) \geq v_i(A_i(\mathbf{b})). \quad (1)$$

We also define notions of truthfulness in the case of randomization: we either have universally truthful CAs, when the mechanism is a probability distribution over deterministic truthful mechanisms, or truthful-in-expectation CAs, when in (1), we use the expected values, over the random coin tosses of the algorithm, of the valuations.

We say that a mechanism A is an α -approximation for CAs with k -minded bidders if for all instances $\mathbf{t} = (v_i, \mathcal{S}_i)_{i=1}^n$, $\sum_{i=1}^n v_i(A_i(\mathbf{t})) \geq \text{OPT}/\alpha$, where OPT denotes the maximum social welfare for instance \mathbf{t} .

We recall that $A_i(t_i, \mathbf{b}_{-i})$ may not belong to the set of demanded sets \mathcal{S}_i ; there can be several sets in \mathcal{S}_i (or none) that are subsets of $A_i(t_i, \mathbf{b}_{-i})$. However, as observed above, the valuation is defined by a set in $\mathcal{S}_i \cup \{\emptyset\}$ which is an inclusion-maximal subset of $A_i(t_i, \mathbf{b}_{-i})$ that maximizes the valuation of agent i . We denote such a set as $\sigma(A_i(t_i, \mathbf{b}_{-i})|t_i)$, i.e., $v_i(A_i(t_i, \mathbf{b}_{-i})) = v_i(\sigma(A_i(t_i, \mathbf{b}_{-i})|t_i))$. In our running example, it can be that for some algorithm A and some \mathbf{b}_{-1} , $A_1(t_1, \mathbf{b}_{-1}) = \{1, 2, 3\} \notin \mathcal{S}_1$. Then, the valuation of $\{1, 2, 3\}$ is defined by $\{1, 2\}$, which is denoted as $\sigma(A_1(t_1, \mathbf{b}_{-1})|t_1)$. Similarly, we define $\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i) \in \mathcal{T}_i \cup \{\emptyset\}$ with respect to $A_i(b_i, \mathbf{b}_{-i})$ and declaration b_i . By the same reasoning, we let $\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i)$ denote the set in $\mathcal{S}_i \cup \{\emptyset\}$ s.t. $v_i(A_i(b_i, \mathbf{b}_{-i})) = v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i))$.

In this work, we focus on *exact* algorithms in the sense of [18], i.e., $A_i(b_i, \mathbf{b}_{-i}) \in \mathcal{T}_i \cup \{\emptyset\}$. Then, since valuations are monotone, $A_i(b_i, \mathbf{b}_{-i}) = \sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)$ and by definition of $\sigma(\cdot)$, for any t_i and b_i in D_i : $\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i) \subseteq A_i(b_i, \mathbf{b}_{-i}) = \sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)$.

In our verification model, each bidder can only declare lower valuations for the set she is awarded. Formally, bidder i with type $t_i = (v_i, \mathcal{S}_i)$ can declare $b_i = (z_i, \mathcal{T}_i)$ iff

$$z_i(A_i(b_i, \mathbf{b}_{-i})) \leq v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i)), \quad (2)$$

when $A_i(b_i, \mathbf{b}_{-i}) \neq \emptyset$. Bidder i evaluates the assigned set $A_i(b_i, \mathbf{b}_{-i}) \in \mathcal{T}_i$ as $\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i) \in \mathcal{S}_i \cup \{\emptyset\}$. Thus, the set $A_i(b_i, \mathbf{b}_{-i})$ can be used to verify *a posteriori* that i has

³We focus on unknown bidders; the discussion naturally adapts to known bidders, where \mathcal{S}_i is known. We can think that \mathcal{S}_i is part of the definition of the private valuation functions v_i . Saying that \mathcal{S}_i is known, we mean that this part of v_i , namely \mathcal{S}_i , is public, anything else about v_i is private.

overbid declaring $z_i(A_i(b_i, \mathbf{b}_{-i})) > v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)) = v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i))$. In the motivating scenario above, the set of cities $A_i(b_i, \mathbf{b}_{-i})$ for which the government assigns licenses to bidder i when declaring b_i , can be used to verify overbidding by simply counting the items available in the stocks of the cities for which licenses are granted to i .

When (2) is not satisfied, the bidder is caught lying by the verification step. We assume that this behavior is very undesirable for the bidder (e.g., in such a case the bidder loses prestige and the possibility to participate in future auctions). This way (1) is satisfied directly when (2) does not hold and truthfulness of an auction is fully captured by (1) holding only for any i , \mathbf{b}_{-i} and $b_i = (z_i, \mathcal{T}_i) \in D_i$ such that (2) is fulfilled. Since we focus on truthful mechanisms with verification and no money, we sometimes avoid to mention that and simply refer to truthful mechanisms/algorithms.

We prove truthfulness by using a variant of the so-called cycle monotonicity technique: Fix algorithm A , bidder i and bids \mathbf{b}_{-i} . The *declaration graph* associated to A has a vertex for each possible declaration in the domain D_i . We add an arc between $a = (z, \mathcal{T})$ and $b = (w, \mathcal{U})$ in D_i whenever a bidder of type a can declare to be of type b obeying (2). Namely, edge (a, b) belongs to the graph iff $z(\sigma(b|a)) \geq w(\sigma(b|b))$.^{4,5} The weight of the edge (a, b) is $z(\sigma(a|a)) - z(\sigma(b|a))$ and encodes the loss that a bidder of type (z, \mathcal{T}) incurs by declaring (w, \mathcal{U}) . The following result relates the weight of edges of the declaration graph to the truthfulness of the algorithm.

PROPOSITION 2.1. *A mechanism A is truthful with verification and without money for CAs with k -minded bidders if and only if each declaration graph associated to A does not have any negative-weight edges.*

3. CHARACTERIZATION OF TRUTHFUL MECHANISMS

We characterize now truthful mechanisms in our setting, for known and unknown bidders. Interestingly, the characterizing property is algorithmic and turns out to be a generalization of the properties used for the design of truthful CAs with money and no verification for single-minded bidders.

3.1 Characterization for Known Bidders

Here, for each k -minded bidder i , we know S_i . The following property generalizes monotonicity of [19] and characterizes truthful auctions without money and with verification.

DEFINITION 3.1. *A mechanism A is k -monotone if for any i , any \mathbf{b}_{-i} , and any $a \in D_i$, if $A_i(a, \mathbf{b}_{-i}) = S$, then for all $b \in D_i$ with $b(S) \geq a(S)$, $b(A_i(b, \mathbf{b}_{-i})) \geq b(S)$.*

THEOREM 3.2. *A mechanism A is truthful with verification and without money for known k -minded bidders if and only if A is k -monotone.*

PROOF. Fix i , \mathbf{b}_{-i} and consider the declaration graph associated to A . Take any edge (b, a) , and let S denote $A_i(a, \mathbf{b}_{-i})$. By definition, the edge exists iff $b(S) \geq a(S)$.

⁴We let $\sigma(b|a)$ be a shorthand for $\sigma(A_i(b, \mathbf{b}_{-i})|a)$ when A , i and \mathbf{b}_{-i} are understood.

⁵From (2), for an edge (a, b) in the declaration graph, we should require that $z(\sigma(b|a)) \geq w(\sigma(b|b))$ only whenever $\sigma(b|b) \neq \emptyset$. But from monotonicity and normalization of valuations, $z(\sigma(b|a)) \geq w(\sigma(b|b))$ also when $\sigma(b|b) = \emptyset$, since $\sigma(b|a) = \emptyset$ and $z(\emptyset) = w(\emptyset) = 0$.

If A is k -monotone, we also have that $b(A_i(b, \mathbf{b}_{-i})) \geq b(S)$, and then the weight $b(A_i(b, \mathbf{b}_{-i})) - b(S)$ of edge (b, a) is non-negative. Vice versa, assume that the weight of (b, a) is non-negative. This means that whenever $b(S) \geq a(S)$, it must be $b(A_i(b, \mathbf{b}_{-i})) \geq b(S)$, and therefore A is k -monotone. The theorem follows from Proposition 2.1. \square

3.2 Characterization for Unknown Bidders

The following property generalizes the property of monotonicity of mechanisms defined by [18] and characterizes truthful auctions without money and with verification.

DEFINITION 3.3. *A mechanism A is k -set monotone if the following holds for any i , any \mathbf{b}_{-i} and any $a = (z, \mathcal{T}) \in D_i$: if $A_i(a, \mathbf{b}_{-i}) = T$ then for all $b = (w, \mathcal{U})$ such that $\sigma(T|b) = U$, $w(U) \geq z(T)$ we have $A_i(b, \mathbf{b}_{-i}) = S$ with $w(S) \geq w(U)$.*

To explain how this notion generalizes [18], we discuss the role of U . In detail, $\sigma(T|b) = U$, in Definition 3.3, should be read as to indicate that bidder i going from declaration a to declaration b , substituted $T \in \mathcal{T}$ with $U \in \mathcal{U}$ and $U \subseteq T$. This is because $\sigma(T|b)$ denotes the set in the collection of sets demanded by a bidder of type b which defines the valuation of T . Specifically, $U \in \mathcal{U}$ is such that $w(U) = w(T)$. (Note that if T belonged to \mathcal{U} , then U would be T itself.) Extending the proof of Theorem 3.2, we show that:

THEOREM 3.4. *A mechanism A is truthful with verification and without money for k -minded bidders if and only if A is k -set monotone.*

PROOF. Fix i , \mathbf{b}_{-i} and consider the declaration graph associated to A . Take any edge $(b = (w, \mathcal{U}), a = (z, \mathcal{T}))$ and let T denote $A_i(a, \mathbf{b}_{-i})$. By definition, the edge exists if and only if $w(U) \geq z(T)$, with $U = \sigma(T|b)$.

Now if the algorithm is k -set monotone, we have that $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$, and the weight $w(A_i(b, \mathbf{b}_{-i})) - w(U)$ of edge (b, a) is non-negative. Vice versa, assume that the weight of (b, a) is non-negative. Hence, whenever $w(U) \geq z(T)$, it must be $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$, and thus A is k -set monotone. The theorem follows from Proposition 2.1. \square

Similarly to [19, 18], k -(set) monotonicity implies the existence of thresholds (a.k.a., critical prices) (for every set). The result in Theorem 3.4 relates to the characterization of truthful CAs with money and no verification (see, e.g., [21, Prop. 9.27]). While the two characterizations look pretty similar, there is an important difference: in the setting with money and no verification, each bidder optimizes her valuation minus the critical price over all her demanded sets; in the setting without money and with verification, each bidder optimizes only her valuation over all her demanded sets among those that are bounded from below by the threshold.

3.3 Implications of Characterizations

A consequence of our results is that a reasonably large class of truthful mechanisms with money can be turned into truthful mechanisms without money but with verification.

3.3.1 Single-Minded versus Multi-Minded Bidders

Our characterization of truthful mechanisms without money for CAs with known and unknown 1-minded bidders is exactly the same as the characterization of truthful mechanisms with money in this setting, see, e.g., [21, pp. 274-275]. Thus, the two classes of truthful mechanisms coincide.

Algorithm 1: Multiplicative price update

- 1 For each good $e \in \mathbf{U}$ do $p_e^1 := p_0$.
 - 2 For each bidder $i = 1, 2, \dots, n$ do
 - 3 $S_i := \arg \max\{v_i(S) : S \in \mathcal{S}_i \text{ s.t. } v_i(S) \geq \sum_{e \in S} p_e^i\}$.
 - 4 Update for each good $e \in S_i$: $p_e^{i+1} := p_e^i \cdot r$.
 - 5 Return $S = (S_1, S_2, \dots, S_n)$.
-

PROPOSITION 3.5. *Any (deterministic) truthful α -approximation mechanism with money for single-minded CAs can be turned into a (deterministic) truthful α -approximation mechanism without money with verification for the same problem, and vice versa. This holds for single-minded CAs with either known or unknown bidders.*

3.3.2 Beyond Combinatorial Auctions

A slight generalization of monotonicity [18] is sufficient to obtain truthful mechanisms with money for problems involving *generalized single-minded bidders* [2]. Intuitively, generalized single-minded bidders have k private numbers in their type: their valuation for a solution is equal to the first of these values or minus infinity, depending on whether the solution asks the agent to “over-perform” on one of the other $k - 1$ parameters, see [2]. By Theorem 3.4, all the truthful mechanisms with money for this quite general type of bidders can be turned into truthful mechanisms without money, when the verification paradigm is justifiable. As a corollary of our characterization, we then have a host of truthful mechanisms without money and with verification for the (multi-objective optimization) problems in [2, 10].

4. UPPER BOUNDS FOR UNKNOWN BIDDERS

In this section, we present our upper bounds on the approximability of CAs with unknown k -minded bidders by truthful mechanisms without money and with verification.

4.1 CAs with Arbitrary Supply of Goods

Next, we consider the more general case where elements in \mathbf{U} are available in b copies each. Note that our characterizations above hold also in this case. We present three polynomial-time algorithms, which are truthful for CAs with unknown bidders: the first is deterministic, the remaining are randomized and give rise to universally truthful CAs.

4.1.1 Deterministic Truthful CAs

We adapt here the overselling multiplicative price update algorithm and its analysis from [17] to our setting. The algorithm is given a parameter $\mu \geq 1$ s.t. $\mu/2 \leq v_{\max} < \mu$. We first assume that μ is known to the mechanism. We then modify our mechanism showing how to truthfully guess v_{\max} .

Algorithm 1 processes the bidders in an arbitrary given order, $i = 1, 2, \dots, n$. The algorithm starts with some relatively small, uniform price $p_0 = \frac{\mu}{4bm}$ of each item. When considering bidder i , the algorithm uses the current prices as defining thresholds and allocates to bidder i a set S_i in her demand set \mathcal{S}_i that has the maximum valuation $v_i(S_i)$ among all her sets with valuations above the thresholds. Then the prices of the elements in the set S_i are increased by a factor r and the next bidder is considered.

Let ℓ_e^i be the number of copies of good $e \in \mathbf{U}$ allocated to all bidders preceding bidder i and $\ell_e^* = \ell_e^{n+1}$ denote the total

Algorithm 2: Modified multiplicative price update

- 1 For each bidder $i \in \{1, 2, \dots, n\}$, let v_{\max}^i be the valuation of i 's most valuable set.
 - 2 Let $j \in \{1, 2, \dots, n\}$ be the bidder with highest value v_{\max}^j (smallest index in case of ties).
 - 3 Let $p_0 = \frac{\mu}{4bm}$, where $\mu = (1 + \epsilon)v_{\max}^j$, for a fixed $0 < \epsilon \ll 1$.
 - 4 For each good $e \in \mathbf{U}$ do $p_e^j := p_0$
 - 5 For any $i = j, 1, 2, 3, 4, \dots, j - 1, j + 1, \dots, n$, let $\text{next}(i)$ be the next number in this order, e.g., $\text{next}(j) = 1$, $\text{next}(1) = 2, \dots, \text{next}(j - 1) = j + 1, \dots$, $\text{next}(n - 1) = n, \text{next}(n) = n + 1$.
 - 6 For each bidder $i = j, 1, 2, \dots, j - 1, j + 1, \dots, n$ do
 - 7 $S_i := \arg \max\{v_i(S) : S \in \mathcal{S}_i \text{ s.t. } v_i(S) \geq \sum_{e \in S} p_e^i\}$.
 - 8 Update for each good $e \in S_i$: $p_e^{\text{next}(i)} := p_e^i \cdot r$.
 - 9 Return $S = (S_1, S_2, \dots, S_{j-1}, S_j, S_{j+1}, \dots, S_n)$.
-

allocation of good e to all bidders. Let, moreover, $p_e^* = p_0 r^{\ell_e^*}$ be good e 's price at the end of the algorithm.

We claim that if p_0 and r are chosen so that $p_0 r^b = \mu$, then the allocation $S = (S_1, \dots, S_n)$ of Algorithm 1 is feasible, i.e., it assigns at most b copies of each good. To show this, fix any good $e \in \mathbf{U}$, and observe that when the b -th copy of e is sold to any bidder, its price becomes $p_0 r^b = \mu > v_{\max}$. Thus, good e alone has a price which is above the maximum valuation of any bidder, and so no further copy will be sold.

Let OPT be the optimal social welfare. We can show that:

THEOREM 4.1. *Algorithm 1 with parameters $p_0 = \frac{\mu}{4bm}$ and $r = (4bm)^{1/b}$ produces a feasible allocation (S_1, \dots, S_n) of social welfare $\sum_{i=1}^n v_i(S_i) \geq \frac{\text{OPT}}{2^{(b(r-1)+1)}} \geq \frac{\text{OPT}}{O(b \cdot (m)^{1/b})}$. Moreover, it is a truthful mechanism without money and with verification for CAs with unknown k -minded bidders.*

We can modify Algorithm 1 and remove the assumption on the knowledge of μ . The modified algorithm is Algorithm 2.

THEOREM 4.2. *Algorithm 2 is a $O(b \cdot (m)^{1/b})$ -approximate truthful mechanism without money and with verification for CAs with unknown k -minded bidders.*

4.1.2 Randomized Truthful CAs

We next discuss how to use Algorithm 2 to obtain randomized universally truthful mechanisms with expected approximation ratios of $O(d^{1/b} \log(bm))$ and $O(m^{1/(b+1)} \log(bm))$.

To this end, observe that if we execute Algorithm 1 with a smaller update factor $r = 2^{1/b}$, the output solution allocates at most sb copies of each good to the bidders, where $s = \log(4bm)$ [17, Lemma 1]. This simply follows from the fact that if sb copies of good $e \in \mathbf{U}$ were sold, then its price would be $p_0 2^{\log(4bm)} = \mu > v_{\max}$. But this infeasible solution is an $O(1)$ -approximation to the optimal feasible solution: plugging $r = 2^{1/b}$ in the approximation ratio of $2^{(b(r-1)+1)}$ in Theorem 4.1 implies an $O(1)$ -approximation (see also [17, Theorem 1]). This idea leads to the following randomized algorithm in [17]: use $r = 2^{1/b}$, explicitly maintain feasibility of the produced solution, and define $q = 1/(2ed^{1/b} \log(4bm))$, where $e \approx 2.718$, as the probability of allocating the best set to a bidder. This algorithm is $O(d^{1/b} \log(4bm))$ -approximate and universally truthful for

Algorithm 3: The greedy algorithm.

- 1 Let l denote the number of different bids, $l = nk$.
 - 2 Let b_1, b_2, \dots, b_l be the non-zero bids and S_1, \dots, S_l be the corresponding sets, ordered such that $b_1 \geq \dots \geq b_l$.
In case of ties between declarations of different bidders consider first the smaller index bidder.
 - 3 For each $j = 1, \dots, l$ let $\beta(j) \in \{1, \dots, n\}$ be the bidder bidding b_j for the set S_j .
 - 4 $\mathcal{P} := \emptyset, \mathcal{B} := \emptyset$.
 - 5 For $i = 1, \dots, l$ do
 - 6 If $\beta(i) \notin \mathcal{B} \wedge S_i \cap S = \emptyset$ for all S in \mathcal{P} then
 - 7 (a) $\mathcal{P} := \mathcal{P} \cup \{S_i\}$, and (b) $\mathcal{B} := \mathcal{B} \cup \beta(i)$.
 - 8 Return \mathcal{P} .
-

CAs with money, cf. [17]. Introducing the same randomization idea into our Algorithm 2, with $r = 2^{1/b}$, we show that:

THEOREM 4.3. *There exists a universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders with an expected approximation ratio of $O(d^{1/b} \cdot \log(bm))$.*

Building on Theorem 4.3, we can also obtain a universally truthful mechanism for demanded sets of unbounded size.

THEOREM 4.4. *There exists a universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders with an expected approximation ratio of $O(m^{1/(b+1)} \cdot \log(bm))$.*

4.2 CAs with Single Supply

We now go back to the case where the goods in U are provided with single supply. In this section, we present three incentive-compatible CAs: the first is deterministic, the remaining two are randomized. Among these three mechanisms, only two run in polynomial time.

4.2.1 Greedy Algorithm

We start with a simple greedy algorithm for CAs where the supply $b = 1$, see Algorithm 3. We note that for goods with arbitrary supply b , the greedy algorithm cannot do better than Algorithm 2 because of the lower bound of \sqrt{m} in [15]. Recall that each bidder $i = 1, 2, \dots, n$ declares (v_i, \mathcal{S}_i) , where \mathcal{S}_i is a collection of k sets bidder i demands and $v_i(S)$ is the valuation of set $S \in \mathcal{S}_i$. Observe that sets S_1, \dots, S_l are all the sets demanded by all bidders (with non-zero bids).

We will use linear programming duality to prove the approximation guarantees of our algorithm. We denote the set family $\mathcal{S} = \cup_{i=1}^n \mathcal{S}_i$, where bidder i demands sets \mathcal{S}_i . For a given set $S \in \mathcal{S}_i$, we denote by $b_i(S)$ the bid of bidder i for that set. Let $[n] = \{1, \dots, n\}$. The LP relaxation of our problem and its corresponding dual linear program are:

$$\begin{aligned}
 & \max \quad \sum_{i=1}^n \sum_{S \in \mathcal{S}_i} b_i(S) x_i(S) \\
 \text{s.t.} \quad & \sum_{i=1}^n \sum_{S: S \in \mathcal{S}_i, e \in S} x_i(S) \leq 1, \quad \forall e \in U \quad (*) \\
 & \sum_{S \in \mathcal{S}_i} x_i(S) \leq 1, \quad \forall i \in [n] \\
 & x_i(S) \geq 0, \quad \forall i \in [n] \forall S \in \mathcal{S}_i \\
 \\
 & \min \quad \sum_{e \in U} y_e + \sum_{i=1}^n z_i \\
 \text{s.t.} \quad & z_i + \sum_{e \in S} y_e \geq b_i(S), \quad \forall i \in [n] \forall S \in \mathcal{S}_i \quad (**) \\
 & z_i, y_e \geq 0, \quad \forall i \in [n] \forall e \in U.
 \end{aligned}$$

THEOREM 4.5. *Algorithm 3 is a $\min\{m, d+1\}$ -approximation algorithm for CAs with unknown k -minded bidders.*

PROOF. Let \mathcal{P} be a solution output by Algorithm 3, and $SAT_{\mathcal{P}} = \cup_{S \in \mathcal{P}} S$. For each set $S \in \mathcal{S}$ with $S \notin \mathcal{P}$, there either is an element $e \in SAT_{\mathcal{P}} \cap S$ which was the *witness* of event $S \notin \mathcal{P}$, or there exists a bidder i and set $S' \in \mathcal{P}$ such that $S', S \in \mathcal{S}_i$. For each set $S \in \mathcal{S} \setminus \mathcal{P}$, we keep in $SAT_{\mathcal{P}}$ one witness for S . If there is more than one witness in $SAT_{\mathcal{P}} \cap S$, we keep in $SAT_{\mathcal{P}}$ the (arbitrary) witness for S that belongs to the set among sets $\{T \in \mathcal{P} : SAT_{\mathcal{P}} \cap S \cap T\}$ that was considered first by the greedy order. We discard the remaining elements from $SAT_{\mathcal{P}}$.

Let us also denote $\mathcal{P}(S) = S \cap SAT_{\mathcal{P}}$ if $S \cap SAT_{\mathcal{P}} \neq \emptyset$ and $\mathcal{P}(S) = S$ if $S \cap SAT_{\mathcal{P}} = \emptyset$.

Observe first that if $m = 1$, then any feasible solution just has a single set assigned to a single bidder and thus the algorithm outputs an optimal solution, as required.

Under the assumption that $m \geq 2$, we will define a dual solution during the execution of Algorithm 3, using the solution \mathcal{P} (\mathcal{P} is needed only for analysis). In line 4 of Algorithm 3 we initialize the duals: $y_e := 0$ for all $e \in U$ and $z_i := 0$ for all $i \in [n]$. We add the following in line 7(a) of Algorithm 3: $y_e := \Delta_e^{S_i}$, for all $e \in \mathcal{P}(S_i)$, where $\Delta_e^{S_i} = \frac{b_{\beta(i)}(S_i)}{|\mathcal{P}(S_i)|}$, for $e \in \mathcal{P}(S_i)$. Note, that for $e \in S_i \setminus SAT_{\mathcal{P}}$ the value of y_e is not updated and remains zero. We also add the following in line 7(a) of Algorithm 3: $z_{\beta(i)} := b_{\beta(i)}(S_i)$.

The following lower bound on the cost of \mathcal{P} is obvious:

$$\sum_{e \in U} y_e \leq \sum_{S_i \in \mathcal{P}} b_{\beta(i)}(S_i). \quad (3)$$

We next show that the scaled solution $(d' \cdot y, z)$ is feasible for the dual linear program, where $d' = \min\{d, m-1\}$. We need to show that **(**)** holds, that is, for each set $S \in \mathcal{S} \cap \mathcal{S}_i$,

$$z_i + d' \sum_{e \in S} y_e \geq b_i(S). \quad (4)$$

Suppose first that $S = S_r \in \mathcal{S} \setminus \mathcal{P}$, and let $\beta(r) = i$. Set S was not included in \mathcal{P} because of two possible reasons: (i) **Case (a):** there is an $e \in SAT_{\mathcal{P}}$ such that $e \in S$, or (ii) **Case (b):** there is a set $S' \in \mathcal{P}$ with $S, S' \in \mathcal{S}_i$.

We first consider Case (a). Then, adding set S to solution \mathcal{P} would violate constraint **(*)**. Let $S'' = S_j \in \mathcal{P}$ be the set in the solution that contains element e and let $h = \beta(j)$.

Recall that $e \in S \cap S''$, thus $\sum_{e' \in S} y_{e'} \geq y_e = \Delta_e^{S''} = \frac{b_h(S'')}{|\mathcal{P}(S'')|} \geq \frac{b_h(S'')}{d} \geq \frac{b_i(S)}{d}$, where the last inequality follows from the greedy selection rule and definition of the witnesses. In the case if $|S| = m$, that is, $S = U$, we obtain that $\sum_{e' \in S} y_{e'} \geq \sum_{e'' \in S''} y_{e''} = \sum_{e'' \in S''} \Delta_{e''}^{S''} = b_h(S'') \geq b_i(S)$, where the last inequality is by the greedy selection rule. Because $m \geq 2$, this proves (4) in Case (a).

We consider now Case (b). Suppose that $S = S_r \in \mathcal{S} \setminus \mathcal{P}$ and there is a set $S' = S_j \in \mathcal{P}$ with $S, S' \in \mathcal{S}_i$. Then we have $i = \beta(j) = \beta(r)$. Observe that when set S' was chosen by Algorithm 3 the dual variable z_i was updated in line 7(a) as follows: $z_i = b_i(S')$. Now, because set S' was considered by the algorithm before set S we have $z_i = b_i(S') \geq b_i(S)$ by the greedy selection rule, which implies (4) in this case.

Claim (4) follows from the definition of z_i if set $S \in \mathcal{S}_i$ is chosen to \mathcal{P} , that is, $S \in \mathcal{P}$. This concludes the proof of (4).

We put all the pieces together. The dual solution $(d' \cdot y, z)$ is feasible for the dual linear program and so by weak

duality $\sum_{i=1}^n z_i + d' \sum_{e \in U} y_e$ is an upper bound on the value of the optimal integral solution to our problem. We have also shown in (3), that $\sum_{e \in U} y_e \leq \sum_{S_i \in \mathcal{P}} b_{\beta(i)}(S_i)$. By letting OPT denote the optimal social welfare, we obtain:

$$\begin{aligned} \text{OPT} &\leq \sum_{i=1}^n z_i + d' \sum_{e \in U} y_e = \sum_{S_i \in \mathcal{P}} z_{\beta(i)} + d' \sum_{e \in U} y_e \\ &\leq \sum_{S_i \in \mathcal{P}} b_{\beta(i)}(S_i) + d' \sum_{S_i \in \mathcal{P}} b_{\beta(i)}(S_i) \\ &= (d' + 1) \sum_{S_i \in \mathcal{P}} b_{\beta(i)}(S_i), \end{aligned}$$

which concludes the proof. \square

THEOREM 4.6. *Algorithm 3 is a truthful mechanism without money and with verification for CAs with unknown k -minded bidders.*

PROOF. Fix i and \mathbf{b}_{-i} . As in Definition 3.3, take two declarations of bidder i , $a = (z, \mathcal{T})$ and $b = (w, \mathcal{U})$ with $w(U) \geq z(T)$, where $T = A_i(a, \mathbf{b}_{-i})$ and $U = \sigma(T|b)$ (here A denotes Algorithm 3). Recall that $U \in \mathcal{U}$ and $U \subseteq T$.

Let \mathcal{S}_a (resp., \mathcal{S}_b) be the set comprised of the sets in declarations of \mathbf{b}_{-i} processed by $A(a, \mathbf{b}_{-i})$ (resp., $A(b, \mathbf{b}_{-i})$) when $z(T)$ (resp., $w(U)$) is considered. Since A grants T to bidder i in the instance (a, \mathbf{b}_{-i}) , it must be the case that $T \cap S = \emptyset$ for all $S \in \mathcal{S}_a$ granted by A . Since $w(U) \geq z(T)$, we have that $\mathcal{S}_b \subseteq \mathcal{S}_a$. Thus, since $U \subseteq T$ then we have that $U \cap S = \emptyset$ for all $S \in \mathcal{S}_b$ granted by the algorithm. Therefore, the only reason for which U might not be granted to i is that A had already granted a set in \mathcal{U} to i , which implies that $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$. Then the algorithm is k -set monotone and the claim follows from Theorem 3.4. \square

[1, Theorem 2] shows a lower bound of $\Omega(m)$ on the approximation ratio of any truthful greedy priority mechanism with money for instances with $d \leq 2$. Nevertheless, Algorithm 3 is truthful without money and with verification and achieves an approximation ratio of 3 for such instances. We next explain the reasons behind this sharp contrast.

PROPOSITION 4.7. *There are no payments that augment Algorithm 3 so to make a truthful mechanism for CAs with k -minded bidders, even in the case of known bidders.*

4.2.2 Randomized Exponential-Time Mechanism

We describe an exponential-time randomized mechanism, or RandExp in brief. Let I be an instance with unknown k -minded bidders, and let I_ℓ , $1 \leq \ell \leq k$, be the subinstance consisting of the elementary bids $(i, S_i^\ell, v_i(S_i^\ell))$, $i \in N$, where S_i^ℓ is the ℓ -th most valuable set demanded by bidder i . Then, RandExp computes the maximum social welfare OPT_ℓ for each subinstance I_ℓ by breaking ties among optimal solutions in a bid-independent way, and outputs the allocation corresponding to OPT_ℓ with probability $1/k$, for each $\ell \in [k]$.

THEOREM 4.8. *RandExp is a k -approximate universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders.*

4.2.3 Randomized Polynomial-Time Mechanism

We conclude with a polynomial-time randomized mechanism, or RandPoly in brief. Let I be an instance of CAs with

unknown k -minded bidders, let v_{\max} be the maximum valuation of some bidder, and let S_{\max} be a set with valuation v_{\max} . Moreover, let I_s be the subinstance that consists of the elementary bids $(i, S, v_i(S))$, $i \in N$, where $|S| \leq \sqrt{m}$. Then, RandPoly either only allocates S_{\max} to the corresponding bidder breaking ties in a bid-independent way with probability $1/2$, or with probability $1/2$, outputs the allocation computed by the Algorithm 3 on the subinstance I_s .

THEOREM 4.9. *RandPoly is a $O(\sqrt{m})$ -approximate universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders.*

5. LOWER BOUNDS FOR KNOWN BIDDERS

We first adapt the proof of [6, Theorem 3.3] and show a lower bound of 2 on the approximation ratio of any deterministic truthful mechanism. We highlight that this lower bound, as well as the lower bounds of Theorems 5.2 and 5.3 below, hold even for exponential-time mechanisms and for simple instances with $n = 2$ bidders and $m = 2$ goods.

THEOREM 5.1. *There are no deterministic truthful mechanisms with approximation ratio better than 2 for CAs with known 2-minded bidders.*

The bound of Theorem 5.1 is tight, since Algorithm 3 gives a 2-approximation for such instances. Theorem 5.1 indicates that our assumption that the bidders do not overbid on their winning sets is less powerful than the use of payments, when we do not take computational issues into account. Furthermore, it shows that already with double-minded bidders, the class of algorithms that can be implemented with money is a strict superset of the class of 2-monotone algorithms.

We next apply Yao's principle and show a lower bound of $5/4$ for randomized universally mechanisms.

THEOREM 5.2. *There are no randomized mechanisms that are universally truthful and have approximation ratio better than $5/4$ for CAs with known 2-minded bidders.*

PROOF. We present a probability distribution over instances with 2 bidders and 2 goods for which the best deterministic truthful mechanism has expected approximation ratio greater than $5/4 - \delta$, for any $\delta > 0$. Let I and I' be two instances on $U = \{a, b\}$. Bidder 1 is interested in $\mathcal{S}_1 = \{\{a, b\}, \{b\}\}$, and bidder 2 is interested in $\mathcal{S}_2 = \{\{a\}\}$. In both, the valuation of bidder 2 is $v_2(\{a\}) = 1$. The valuation of bidder 1 is $v_1(\{a, b\}) = 2$ and $v_1(\{b\}) = 0$ in I , and $v'_1(\{a, b\}) = 2$ and $v'_1(\{b\}) = 2 - \delta$ in I' . Each instance occurs with probability $1/2$, and the expected maximum social welfare is $(5 - \delta)/2$. Let algorithm A , applied to instance I , allocate $\{a, b\}$ to bidder 1 and \emptyset to bidder 2. Then, by Theorem 3.2, since A is a deterministic truthful mechanism, when applied to instance I' , it must allocate $\{a, b\}$ to bidder 1 and \emptyset to bidder 2. Therefore, the expected social welfare of A is 2, and its expected approximation ratio is $(5 - \delta)/4 > 5/4 - \delta$. If A , applied to I , does not allocate $\{a, b\}$ to bidder 1, its expected social welfare is at most $(4 - \delta)/2$, and its approximation ratio is $(5 - \delta)/(4 - \delta) > 5/4 - \delta$, a contradiction. \square

We conclude with a weaker lower of 1.09 for the larger class of randomized truthful-in-expectation mechanisms.

THEOREM 5.3. *There are no randomized mechanisms that are truthful in expectation and have approximation ratio less than 1.09 for CAs with known 2-minded bidders.*

5.1 Priority Mechanisms

Next, we consider mechanisms that operate according to the priority framework (see e.g., [1]). We start with any priority mechanism A that takes as input (and make irrevocable decisions about) elementary bids. Namely, A operates on a sequence of triples $(i, S, v_i(S))$, where i is the bidder, S is one of i 's demanded sets, and $v_i(S)$ is i 's valuation for S .

THEOREM 5.4. *Let A be a truthful priority mechanism with verification and no money for CAs with known k -minded bidders. If A processes elementary bids, the approximation ratio of A is greater than $(1 - \delta)d$, for any $\delta > 0$.*

The proof of Theorem 5.4 adapts that of [1, Theorem 3]. With a minor change in the proof, Theorem 5.4 applies to the special case of 2-minded bidders. Thus, exploiting instances with $d = m$, we obtain a lower bound of $(1 - \delta)m$, on the approximation ratio of any truthful priority mechanism for known 2-minded bidders that processes elementary bids.

Next, we consider any priority mechanism A that takes as input k -minded bidders, i.e., A operates on a sequence of pairs (i, v_i) , where v_i is the valuation function of bidder i . Such a priority mechanism A is potentially more powerful than a priority mechanism processing elementary bids, since when A decides about the set allocated to each bidder i , it has full information about i 's valuation function. The proof of the following adapts the proof of [1, Theorem 4].

THEOREM 5.5. *Let A be a truthful priority mechanism with verification and no money for CAs with known 2-minded bidders. If A processes bidders, the approximation ratio of A is greater than $(1 - \delta)m/2$, for any $\delta > 0$.*

5.2 Discussion

A step that seems necessary for $O(\sqrt{m})$ -approximation for CAs is that the algorithm compares the social welfare and chooses the best of two extreme solutions: the most valuable set demanded by some bidder and a solution consisting of many small sets with a large total valuation. Otherwise, we cannot achieve an approximation ratio of $o(m)$ even for the simple case where bidder 1 is double-minded for $\mathbf{U} = \{a_1, \dots, a_m\}$ with valuation $x \in \{1 + \varepsilon, m^2\}$ and for the good a_1 with valuation 1, and each bidder i , $2 \leq i \leq m$, is single-minded for the good a_i with valuation 1. In fact, this is one of the restrictions of priority algorithms exploited in the proofs of the lower bounds of $\Omega(m)$ above.

On the other hand, comparing the social welfare of these two extreme solutions is also sufficient for an $O(\sqrt{m})$ -approximation, in the sense that taking the best of (i) the most valuable set demanded by some bidder, and (ii) the solution of Algorithm 3, if we only allocate sets of cardinality at most \sqrt{m} , is an $O(\sqrt{m})$ -approximation (see Theorem 4.9).

For CAs without money, it seems virtually impossible to let a deterministic mechanism truthfully implement a comparison between the social welfare of those extreme solutions. This is because the only way for a deterministic mechanism to make sure that the bidder with the maximum valuation does not lie about it is to allocate her most valuable set to her, so that verification applies to this particular bid (see also how Algorithm 2 learns about v_{\max}). But this leads to an approximation ratio of $\Omega(m)$. In fact, this is the main obstacle towards a deterministic truthful $O(\sqrt{m})$ -approximate mechanism for CAs with k -minded bidders. So, we conjecture that there is a lower bound of $\Omega(m)$ on the approximation ratio of deterministic truthful mechanisms.

6. REFERENCES

- [1] A. Borodin and B. Lucier. On the limitations of greedy mechanism design for truthful combinatorial auctions. *ICALP*, pp. 90-101, 2010.
- [2] P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanism design. *STOC*, pp. 39-48, 2005.
- [3] G. Carroll. When are local incentive constraints sufficient? *Econometrica*, 80(2): 661-686, 2012.
- [4] G. Celik. Mechanism design with weaker incentive compatibility constraints. *Games and Economic Behavior*, 56(1): 37-44, 2006.
- [5] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. *APPROX*, pp. 89-103, 2007.
- [6] S. Dughmi and A. Ghosh. Truthful assignment without money. *EC*, 2010.
- [7] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. *SODA*, 2004.
- [8] D. Fotakis, P. Krysta, and C. Ventre. Combinatorial auctions without money. *CoRR*, abs/1310.0177, 2013.
- [9] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple Facility Location games. *Theor. Comp. Sc.*, 472: 90-103, 2013.
- [10] F. Grandoni, P. Krysta, S. Leonardi, and C. Ventre. Utilitarian mechanism design for multi-objective optimization. *SODA*, pp. 573-584, 2010.
- [11] J. R. Green and J. Laffont. Partially Verifiable Information and Mechanism Design. *The Review of Economic Studies*, 53: 447-456, 1986.
- [12] E. Hazan, S. Safra, and O. Schwartz. On the complexity of approximating -set packing. *Computational Complexity*, 15(1): 20-39, 2006.
- [13] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in CAs. *Games and Economic Behavior*, 47: 104-123, 2004.
- [14] E. Koutsoupias. Scheduling without payments. *SAGT*, pp. 143-153, 2011.
- [15] P. Krysta. Greedy approximation via duality for packing, combinatorial auctions and routing. *MFCS*, pp. 615-627, 2005.
- [16] P. Krysta and C. Ventre. Combinatorial auctions with verification are tractable. *ESA*, pp. 39-50, 2010.
- [17] P. Krysta and B. Vöcking. Online mechanism design (randomized rounding on the fly). *ICALP*, pp. 636-647, 2012.
- [18] D. J. Lehmann, L. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5): 577-602, 2002.
- [19] A. Mu'Alem and N. Nisan. Truthful approximation mechanisms for restricted CAs. *AAAI*: 379-384, 2002.
- [20] N. Nisan. The communication complexity of approximate set packing and covering. *ICALP*, 2002.
- [21] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. 2007.
- [22] P. Penna and C. Ventre. Collusion-resistant mechanisms with verification yielding optimal solutions. *ESA*, pp. 708-719, 2008.
- [23] P. Penna and C. Ventre. Optimal collusion-resistant mechanisms with verification. *EC*: 147-156, 2009.
- [24] A. Procaccia and M. Tennenholtz. Approximate mechanism design without money. *EC*: 177-186, 2009.