

Combinatorial Auctions Without Money

Dimitris Fotakis¹ · Piotr Krysta² ·
Carmine Ventre³

Received: 30 April 2015 / Accepted: 19 December 2015 / Published online: 29 December 2015
© Springer Science+Business Media New York 2015

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. Additionally, truthfulness of CAs is poorly understood already in the case in which bidders happen to be interested in only two different sets of goods. 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

Dimitris Fotakis is supported by the project AlgoNow, co-financed by the European Union (European Social Fund—ESF) and Greek national funds, through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF)—Research Funding Program: THALES, investing in knowledge society through the European Social Fund. Piotr Krysta is supported by EPSRC Grant EP/K01000X/1. Carmine Ventre is supported by EPSRC Grant EP/M018113/1.

✉ Carmine Ventre
c.ventre@tees.ac.uk

Dimitris Fotakis
fotakis@cs.ntua.gr

Piotr Krysta
pkrysta@liverpool.ac.uk

¹ National Technical University of Athens, Athens, Greece

² University of Liverpool, Liverpool, UK

³ Teesside University, Middlesbrough, UK

verification and without money. We prove a characterization of truthful mechanisms, which makes an interesting parallel with the well-understood case of CAs with money for single-minded bidders. We then give a host of upper bounds on the approximation ratio obtained by either deterministic or randomized truthful mechanisms when the sets and valuations are private knowledge of the bidders. (Most of these mechanisms run in polynomial time and return solutions with (nearly) best possible approximation guarantees.) We complement these positive results with a number of lower bounds (some of which are essentially tight) that hold in the easier case of public sets. We thus provide an almost complete picture of truthfully approximating CAs in this general setting with multi-dimensional bidders.

Keywords Combinatorial auctions · Algorithmic Mechanism Design · Approximate mechanism design without money · Mechanisms with verification

1 Introduction

Algorithmic Mechanism Design has as main scope the realignment of the objective of the designer with the selfish interests of the agents involved in the computation. Since the Internet, as the principal computing platform nowadays, is perhaps the main motivation to study problems in which these objectives are different, one would expect truthful mechanisms to have concrete and widespread practical applications. However, one of the principal obstacles to this is the assumption that the mechanisms use monetary transfers. On one hand, money may provoke (unreasonably) large payments [9]; on the other hand, while money might be reasonable in some applications, such as sponsored search auctions, 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 is 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, for short), the paradigmatic problem in Algorithmic Mechanism Design. In a combinatorial auction 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 ($v_i(\emptyset)$ is normalized to be 0). 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 $\sum_{i=1}^n v_i(S_i)$ —the *social welfare*—is maximized. For this problem, 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 and deal with budgeted bidders (a case which is generally hard to handle in presence of money).

In this paper, we focus on k -minded bidders, i.e., bidders are interested in obtaining one out of a collection of k subsets of U . In this general setting, we want to study the feasibility of designing truthful CAs without money, returning (ideally, in polyno-

mial time) 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 [8]. We therefore focus on the model of CAs with verification, introduced in [18]. In this model, which is motivated by a number of real-life applications and has also been considered by economists [6], bidders do not overbid their valuations on the set that they are awarded. 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.

The model of CAs with verification is perhaps best illustrated by means of the following motivating scenario, discussed first in [18]. Consider a government (auctioneer) auctioning business licenses for a set U of cities under its administration. A company (bidder) wants to get a business 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 in literature in Economics [6], a simple inspection on the stock consistency implies that bidders cannot overbid their profits: the concealment of existing product items in stock is costless but disclosure of unavailable ones is prohibitively costly. The assumption is verification a posteriori²: the inspection is carried on for the solutions 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. A formal definition of the model of CAs with verification and without money can be found in Sect. 2.

1.1 Our Contribution

In this model, we firstly give a complete characterization of algorithms that are incentive-compatible in both the cases in which the collections of k sets, each bidder is interested in, are public (also referred to, as known bidders) and private (also known as, unknown bidders); valuations are always assumed to be private. We prove that truthfulness is characterized in this context in terms of k -monotone algorithms: in the case of known bidders, if a bidder is awarded a set S and augments her declaration for S then a k -monotone algorithm must, in this new instance, grant her a set in her collection which is not worse than S (i.e., a set with a valuation not smaller than her valuation for S). (This generalizes neatly to the case of unknown bidders.)

¹ Note that bidders will sell products already in stock (i.e., no production costs are involved as they have been sustained before the auction is run). This is conceivable when a government runs an auction for urgent needs (e.g., salt provision for icy roads or vaccines for pandemic diseases).

² A 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.

There are two important facts we wish to emphasize about our characterizations. First, their significance stems from the fact that the corresponding problem of characterizing truthfulness for CAs with money and k -minded bidders is poorly understood: this is a long-standing open problem already for $k = 2$, see, e.g., [26, Chapter 12]. Second, it is pretty easy to see that these notions generalize the properties of monotonicity shown to characterize truthfulness with money for single-minded bidders in [20, 23] for known and unknown bidders, respectively. More generally, these properties of monotonicity are also proved to be sufficient to get truthful mechanisms for so-called generalized single-minded bidders [3]. This is an interesting development as, to the best of our knowledge, 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 then their approximation guarantees) remain unchanged. Thus, in light of our results, previously known algorithms presented in, e.g., [3, 12, 20] assume a double relevance: they are truthful not only when money can be used, but also in absence of money when verification can be implemented. This equivalence gives also a strong motivation for our model. Naturally, the picture for the multi-dimensional case of $k > 1$ is more blurry since, as we mention above, truthfulness with money is not well understood yet in these cases.

Armed with the characterization of truthfulness, we provide a number of upper and lower bounds on the approximation guarantee to the optimal social welfare of truthful CAs without money and with verification. The upper bounds hold for the harder case of unknown bidders. We give an upper bound of $O(b\sqrt[b]{m})$ in the case in which each good in U has a supply b . This algorithm is deterministic, runs in polynomial time and adapts an idea of multiplicative update of good prices by [19]. Following similar ideas, we also obtain randomized universally truthful mechanisms with approximation ratios of $O(d^{1/b} \cdot \log(bm))$ and $O(m^{1/(b+1)} \cdot \log(bm))$, 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. We also give two simple randomized universally truthful CAs without money: 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—cf. Table 1. We also would like to note that all, but the k -approximate, upper bounds given can be obtained in the setting in which bidders’ declare so-called demand oracles, see, e.g., [26, Chapter 11].

We complete this study by proving a host of lower bounds on the approximation guarantee of truthful CAs without money for known bidders, without any computational assumption. (Note that the class of incentive-compatible algorithms for known bidders is larger than the class for unknown bidders.) 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. Additionally, stronger lower bounds are proved for deterministic truthful mechanisms that use priority algorithms [1]. These

Table 1 Our upper bounds compared with the computational hardness to approximate the problem

Mechanism type	Upper bound (unknown k -minded bidders)	Computational lower bounds [14, 20, 24]
Deterministic	$O(b \cdot m^{1/b})$	
Univ. truthful	$O(d^{1/b} \cdot \log(bm))$	
Univ. truthful	$O(m^{1/(b+1)} \cdot \log(bm))$	$O(d/\log d)$
$b = 1$, deterministic	$\min\{m, d + 1\}$	$m^{1/2-\epsilon}$
$b = 1$, univ. truthful	k (exp. time)	
$b = 1$, univ. truthful	$O(\sqrt{m})$	

Table 2 Our lower bounds at a glance

Mechanism type	Unconditional lower bounds (known k -minded bidders)
Deterministic	2
Univ. truthful	5/4
Truthful in expectation	1.09
Priority (bids)	$(1 - \delta)d$, any $\delta > 0$
Priority (bidders)	$(1 - \delta)m/2$, any $\delta > 0$

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. (Note that our greedy mechanism falls in the category of non-adaptive priority algorithms since it process bids as items, which are ordered at the beginning.) 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. A summary of our lower bounds is presented in Table 2.

Our bounds give a rather surprising picture of the relative power of verification versus money, thus suggesting that the two models are somehow incomparable. 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 [7]. (However, it is worth mentioning that the latter mechanism does not guarantee the approximation ratio since there is an error probability of $O(\log m/\sqrt{m})$ which cannot be reduced by, e.g., repeating the auction or otherwise truthfulness would be lost.) On the other hand, because of our lower bounds, we know that it is not possible to implement the optimal outcome without money; while, if we have exponential computational time, we can truthfully implement the optimal solution using VCG payments. However, if we restrict to polynomial-time mechanisms, then we have a deterministic truthful $\min\{m, d + 1\}$ -approximation mechanism without money, based on the aforementioned greedy algorithm; with money, instead, it is not known how to obtain any polynomial-time deterministic truthful mechanism with an approximation ratio better than the $O(m/\sqrt{\log m})$ -approximation given in [15]. Moreover, [1, Theorem 2] proved a lower bound of $\Omega(m)$ on the approximation ratio of any truthful greedy mechanism with money for instances with demanded sets of cardinality at most 2. Our greedy

mechanism achieves an approximation ratio of 3 for such instances, which implies that this lower bound 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. In this regard, we offer an interesting conjecture in Sect. 5.5.

1.2 Related Work

CAs as an optimization problem (without strategic consideration) is known to be 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 [14, 20, 24]. As a consequence, 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. For single-minded domains, a host of truthful CAs have been designed (see, e.g., [3, 20, 23]). A more complete picture of what is known for truthful CAs under different restrictions of bidders' domains can be found in Figure 11.2 of [26].

The authors of [18], instead of restricting the domains of the bidders, proposed to restrict the way bidders lie. We are adopting 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., [25, 28, 29] and references therein). Mechanism design problems where players have restrictions on the way of lying are also considered in theoretical economics. We next discuss some of the work more relevant to this paper. Green and Laffont [13] define and motivate a model of partial verification wherein bidders can only report bids from a type-dependent set of allowed messages; they characterize bidding domains for which the Revelation Principle holds in presence of this notion of restricted bidding. This model has been further studied by Singh and Wittman [32] and later extended in [4] to allow probabilistic verification of bids outside the set of allowed messages. The economic model that is closest to ours is the one studied in [6]; therein verification is supposed to take place for every outcome and not just for the implemented solution and is therefore stronger and less realistic than ours. Another related line of work tries to establish when a subset of incentive-compatibility constraints is sufficient to obtain full incentive-compatibility. Moore [22] considers a single good, single buyer optimal auction design and studies conditions under which no-overbidding constraints would also imply the full incentive compatibility of the underlying auction. Other papers studying this kind of questions are [5, 31]. In particular, the results in [5, 11] (and to some extent in [4]) imply that one has to focus only on “one-sided” verification, for otherwise a mechanism is truthful if and only if it satisfies a subset of incentive-compatibility constraints.

Our work fits in the framework of approximate mechanism design without money, initiated by [30]. 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 algorithms. Approximate mechanisms without money have been obtained for various

problems, among them, for locating one or two facilities in metric spaces (see e.g., [21,30]). Due to the apparent difficulty of truthfully locating three or more facilities with a reasonable approximation guarantee, notions conceptually similar to our notion of verification have been proposed [10,27]. Koutsoupias [16] considers truthful mechanisms without money, for scheduling selfish machines whose execution times can be (strongly) verified. The authors of [8] consider the design of mechanisms without money for, what they call, the Generalized Assignment problem: 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 modeled via maximum weight bipartite matching and the latter can be cast as a special case of CAs with demanded sets of cardinality 1; then [8, Algorithm 1] can be regarded as a special case of our greedy algorithm.

2 Model and Preliminaries

In a combinatorial auction we have a set \mathbf{U} of m goods and n agents, a.k.a. bidders. Each k -minded 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 \mathbf{U} , k being the size of \mathcal{S}_i . The valuation function maps subsets of goods to nonnegative real numbers ($v_i(\emptyset)$ is normalized to be 0). Agents' valuations are 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 \mathbf{U} such that $\sum_{i=1}^n v_i(S_i)$ –the *social welfare*– is maximized. As an example consider $\mathbf{U} = \{1, 2, 3\}$ 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

$$v_i(S) = \begin{cases} \max_{S' \in \mathcal{S}_i: S \supseteq S'} \{v_i(S')\} & \text{if } \exists S' \in \mathcal{S}_i \wedge S \supseteq S', \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Accordingly, 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 $v_i(S) = 0$ then we say that \emptyset defines it. So in the example above $v_1(\{1, 2, 3\})$ is defined by $\{1, 2\}$.

Throughout the paper we assume that bidders are interested in sets of cardinality at most $d \in \mathbb{N}$, i.e., $d = \max\{|S| : \exists i \text{ s.t. } S \in \mathcal{S}_i \wedge v_i(S) > 0\}$.

Assume that the sets $S \in \mathcal{S}_i$ and the values $v_i(S)$ are private knowledge of the bidders. Then, we want to design an *allocation algorithm (auction)* that for a given input of bids from the bidders, outputs a feasible assignment (i.e., at most one of the requested sets is allocated to each bidder, and allocated sets are pair-wise disjoint). The auction should guarantee that no bidder has an incentive to misreport her preferences and maximize the social welfare (i.e., the sum of the valuations of the winning bidders).

More formally, we let \mathcal{T}_i be a set of k non-empty subsets of \mathbf{U} and let z_i be the corresponding valuation function of agent i , i.e., $z_i : \mathcal{T}_i \rightarrow \mathbb{R}^+$. 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 agent i . We also let D_i denote the set of all the possible declarations of agent i and call D_i the *declaration domain* of bidder i . Fix the declarations \mathbf{b}_{-i} of all the agents but i . For any declaration $b_i = (z_i, \mathcal{T}_i)$ in D_i , we let $A_i(b_i, \mathbf{b}_{-i})$ be the set that an auction A on input $\mathbf{b} = (b_i, \mathbf{b}_{-i})$ allocates to bidder i . If no set is allocated to i then we naturally

set $A_i(b_i, \mathbf{b}_{-i}) = \emptyset$. Observe that, according to (1), $v_i(\emptyset) = 0$. We say that A is a truthful auction without money if the following holds for any $i, b_i \in D_i$ and \mathbf{b}_{-i} :

$$v_i(A_i(t_i, \mathbf{b}_{-i})) \geq v_i(A_i(\mathbf{b})). \tag{2}$$

We also define notions of truthfulness in the case of randomization: we either have universally truthful CAs, in which case the mechanism is a probability distribution over deterministic truthful mechanisms, or truthful-in-expectation CAs, where in (2) we use the expected values, over the random coin tosses of the algorithm, of the valuations. We also say that a mechanism A is an α -approximation for CAs with k -minded bidders if for all $\mathbf{t} = (v_i, \mathcal{S}_i)_{i=1}^n, \sum_{i=1}^n v_i(A_i(\mathbf{t})) \geq \text{OPT}/\alpha$, OPT being the value of a solution with maximum social welfare for the instance \mathbf{t} .

Recall that $A_i(t_i, \mathbf{b}_{-i})$ may not belong to the set of demanded sets \mathcal{S}_i . In particular, 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 (cf. (1)), the valuation is defined by a set in $\mathcal{S}_i \cup \{\emptyset\}$ which is an inclusion-maximal subset of set $A_i(t_i, \mathbf{b}_{-i})$ 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 above, it can be for some algorithm A and some \mathbf{b}_{-1} , that $A_1(t_1, \mathbf{b}_{-1}) = \{1, 2, 3\} \notin \mathcal{S}_1$ whose valuation is defined as observed above by $\{1, 2\}$; the set $\{1, 2\}$ is denoted as $\sigma(A_1(t_1, \mathbf{b}_{-1})|t_1)$. (Similarly, we define $\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i) \in \mathcal{T}_i \cup \{\emptyset\}$ w.r.t. $A_i(b_i, \mathbf{b}_{-i})$ and declaration b_i .) Following the same reasoning, we let $\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i)$ denote the set in $\mathcal{S}_i \cup \{\emptyset\}$ such that $v_i(A_i(b_i, \mathbf{b}_{-i})) = v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i))$.

We focus on *exact* algorithms³ in the sense of [20]. This means that $A_i(b_i, \mathbf{b}_{-i}) \in \mathcal{T}_i \cup \{\emptyset\}$. This implies, by monotonicity of the valuations, that $A_i(b_i, \mathbf{b}_{-i}) = \sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)$ and then the definition of $\sigma(\cdot)$ yields the following for any $t_i, b_i \in D_i$:

$$\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i) \subseteq A_i(b_i, \mathbf{b}_{-i}) = \sigma(A_i(b_i, \mathbf{b}_{-i})|b_i). \tag{3}$$

Verification model for CAs In the verification model each bidder can only declare lower valuations for the set she is awarded. More formally, bidder i whose type is $t_i = (v_i, \mathcal{S}_i)$ can declare a type $b_i = (z_i, \mathcal{T}_i)$ if and only if whenever $\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i) \neq \emptyset$:

$$z_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)) \leq v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i)). \tag{4}$$

In particular, bidder i evaluates the assigned set $\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i) \in \mathcal{T}_i$ as $\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i) \in \mathcal{S}_i \cup \{\emptyset\}$, i.e., $v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i)) = v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i))$. Thus the set $\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)$ can be used to verify *a posteriori* that bidder i has overbid declaring $z_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)) > v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|b_i)) = v_i(\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i))$. To be more concrete, consider the motivating scenario for CAs with verification above. The set of cities $\sigma_i(A(b_i, \mathbf{b}_{-i})|b_i)$ for which the government assigns licenses to bidder i when declaring b_i , can be used *a posteriori* to verify overbidding by simply counting the product items available in the stock of the cities for which licenses were granted to bidder i .

³ An algorithm is exact if, to each bidder, either only one of the declared sets is awarded or none.

When (4) is not satisfied then the bidder is caught lying by the verification step. We assume that this behavior is very undesirable for the bidder (e.g., for simplicity we can assume that in such a case the bidder loses prestige and the possibility to participate in the future auctions). This way (2) is satisfied directly when (4) does not hold (as in such a case a lying bidder would have an infinitely bad utility because of the assumption above). Thus in our model, truthfulness with verification and without money of an auction is fully captured by (2) holding only for any i , \mathbf{b}_{-i} and $b_i = (z_i, T_i) \in D_i$ such that (4) is fulfilled. Since our main focus is on this class of truthful mechanisms with verification and no money, we sometimes avoid mentioning that and simply refer to truthful mechanisms/algorithms.

A graph-theoretic approach The technique we will use to derive truthful auctions for multi-minded XOR bidders is a straightforward variation of the so-called cycle monotonicity technique. Consider an algorithm A . We will set up a weighted graph for each bidder i depending on A , bidder domain D_i and the declaration \mathbf{b}_{-i} of all the bidders but i in which the non-existence of negative-weight edges guarantees the truthfulness of the algorithm. This is a well known technique. More formally, fix algorithm A , bidder i and declarations \mathbf{b}_{-i} . The *declaration graph* associated to algorithm A has a vertex for each possible declaration in the domain D_i . We add an arc between $a = (z, T)$ and $b = (w, U)$ in D_i whenever a bidder of type a can declare to be of type b obeying (4). Following the definition of the verification setting, edge (a, b) belongs to the graph if and only if $z(\sigma(b|a)) \geq w(\sigma(b|b))$.^{4 5} The weight of the edge (a, b) is defined as $z(\sigma(a|a)) - z(\sigma(b|a))$ and thus encodes the loss that a bidder whose type is (z, T) incurs by declaring (w, U) . The following result (whose proof is straightforward) relates the weight of edges of the declaration graph to the truthfulness of the algorithm.

Proposition 1 *A is a truthful auction with verification without money for CAs with k -minded bidders if and only if each declaration graph associated to algorithm A does not have negative-weight edges.*

In the case of mechanisms without verification, the graph above is complete. Such a graph can be used to check whether algorithms can be augmented with payments so to ensure truthfulness, both in the scenario with verification and without. Incentive-compatibility of algorithms is known to coincide with the case in which each graph has no negative-weight cycles [33]. We will use this fact to show that certain algorithms cannot be made truthful with money.

Known versus Unknown k -minded bidders In the discussion above, we consider the case in which the collection of k sets, each bidder is interested in, is private knowledge. In this case, we refer to the problem of designing truthful auctions that maximize the social welfare as CAs with *unknown k -minded bidders* (or, simply, *unknown bidders*). An easier scenario is the setting in which the sets are public knowledge and bidders

⁴ To ease our notation we let $\sigma(b|a)$ be a shorthand for $\sigma(A_i(b, \mathbf{b}_{-i})|a)$ when the algorithm, the bidder i and declarations \mathbf{b}_{-i} are clear from the context as in this case.

⁵ Strictly speaking for an edge (a, b) in the graph, we should require that $z(\sigma(b|a)) \geq w(\sigma(b|b))$ only whenever $\sigma(b|b) \neq \emptyset$ as this set would be needed to verify. However, because of the monotonicity and normalization of valuations, $z(\sigma(b|a)) \geq w(\sigma(b|b))$ holds also whenever $\sigma(b|b) = \emptyset$, since $\sigma(b|a) = \emptyset$ and $z(\emptyset) = w(\emptyset) = 0$.

are only strategic about their valuations. In this case, we instead talk about CAs with *known k -minded bidders* (or, simply, known bidders). Our upper bounds hold for the more general case of unknown bidders, while the lower bounds apply to the larger class of mechanisms truthful for known bidders.

3 Characterization of Truthful Mechanisms

In this section we characterize the algorithms that are truthful in our setting, in both the scenarios of known and unknown bidders. Interestingly, the characterizing property is algorithmic only 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

In this case, for each k -minded bidder i we know S_i . The following property generalizes monotonicity of [23] and characterizes truthful auctions without money and with verification.

Definition 1 An algorithm A is *k -monotone* if the following holds for any i , any \mathbf{b}_{-i} , any $a \in D_i$: if $A_i(a, \mathbf{b}_{-i}) = S$ then for all $b \in D_i$ such that $b(S) \geq a(S)$ it holds $b(A_i(b, \mathbf{b}_{-i})) \geq b(S)$.

To gain some intuition, we give an example explaining what this definition means for known single-minded bidders, see, e.g., [23]. If bidder i is single-minded, then her declaration $a_i = (z_i, T_i) \in D_i$ contains only a single subset, that is, $T_i = \{T\}$, for some $T \subseteq U$, and then, for any subset $S \subseteq U$, (1) reduces to

$$z_i(S) = \begin{cases} z_i(T) & \text{if } S \supseteq T, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

Now, suppose that on declaration a_i the mechanism assigns to bidder i set S , that is, $A_i(a_i, \mathbf{b}_{-i}) = S$ and bidder i changes her declaration to $b_i = (w_i, T_i)$ (note that $T_i = \{T\}$ remains the same because the bidder is known single-minded and w_i fulfills (5)). Applying Definition 1 of 1-monotonicity, if $b_i(S) \geq a_i(S)$, which is equivalent to $w_i(S) \geq z_i(S)$, then $b_i(A_i(b_i, \mathbf{b}_{-i})) \geq b_i(S)$. The only non-trivial case is when $S = T$ (recall that mechanism A_i is exact) and $a_i(S) = z_i(T) > 0$. But then $b_i(A_i(b_i, \mathbf{b}_{-i})) > 0$ and by exactness of A_i , we have that $b_i(A_i(b_i, \mathbf{b}_{-i})) = T$. This means that Definition 1 is the same in case of known single-minded bidders as the definition of monotonicity in [23].

Theorem 1 An algorithm A is truthful without money and with verification 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 algorithm A . Take any edge of the graph (b, a) and let S denote $A_i(a, \mathbf{b}_{-i})$. By definition, the edge exists if and only if $b(S) \geq a(S)$.

Now if the algorithm is k -monotone then 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 $b(S) \geq a(S)$ then it must be $b(A_i(b, \mathbf{b}_{-i})) \geq b(S)$ and therefore A is k -monotone. The theorem follows from Proposition 1. \square

Similarly to [23], k -monotonicity implies the existence of thresholds (critical values/prices). Towards this end, it is important to consider the sets in S_i in decreasing order of (true) valuations. Accordingly, we denote $S_i = \{S_i^1, \dots, S_i^k\}$, with $v_i(S_i^j) > v_i(S_i^l)$ if and only if $j < l$.

Lemma 1 *An algorithm A is k -monotone if and only if for any i , any \mathbf{b}_{-i} , any t_i there exist k threshold values $\Theta_i^1(\mathbf{b}_{-i}), \dots, \Theta_i^k(\mathbf{b}_{-i})$ such that: if $b_i(S_i^j) > \Theta_i^j(\mathbf{b}_{-i})$ and $b_i(S_i^\ell) < \Theta_i^\ell(\mathbf{b}_{-i})$, for all $\ell < j$ then $\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i) = S_i^j$. Moreover, if $b_i(S_i^\ell) < \Theta_i^\ell(\mathbf{b}_{-i})$, for all $\ell \in [k]$ then $\sigma(A_i(b_i, \mathbf{b}_{-i})|t_i) = \emptyset$.*

Proof Let us first show that the existence of the thresholds for an algorithm A implies that A is k -monotone. Fix i, \mathbf{b}_{-i} and t_i ; let a be a declaration in D_i such that $a(S_i^j) > \Theta_i^j(\mathbf{b}_{-i})$ for some $j \in [k]$ and $a(S_i^\ell) < \Theta_i^\ell(\mathbf{b}_{-i})$ for $\ell < j$. Then $A_i(a, \mathbf{b}_{-i}) = S_i^j$. Now take $b \in D_i$ such that $b(S_i^j) \geq a(S_i^j)$. No matter what the declaration $b(S_i^h)$ is, for $h \neq j$, by definition of thresholds we have that $A_i(b, \mathbf{b}_{-i}) = S_i^l, l \leq j$.

For the opposite direction, fix i, \mathbf{b}_{-i} and t_i ; let a be a declaration in D_i such that $A_i(a, \mathbf{b}_{-i}) = S_i^j$. We prove by induction on j that there exists a threshold $\Theta_i^j(\mathbf{b}_{-i})$ as in the statement. For the base case consider $j = 1$. It is straightforward to see that the threshold $\Theta_i^1(\mathbf{b}_{-i})$ exists [20]. Now assume that there are thresholds $\Theta_i^1(\mathbf{b}_{-i}), \dots, \Theta_i^{j-1}(\mathbf{b}_{-i})$ and take $b \in D_i$ to be such that $b(S_i^j) \geq a(S_i^j)$ and $b(S_i^h) < \Theta_i^h(\mathbf{b}_{-i})$ for $h < j$. By definition of $\Theta_i^1(\mathbf{b}_{-i}), \dots, \Theta_i^{j-1}(\mathbf{b}_{-i})$ and that of k -monotonicity, $\sigma(A_i(b, \mathbf{b}_{-i})|t_i) = S_i^j$. We can then conclude that there exists a threshold $\Theta_i^j(\mathbf{b}_{-i})$ as well. \square

The lemma above assumes that bidders have k different valuations for each of their minds. This is a rather nonrestrictive way to model CAs for k -minded bidders. In the more general case in which bidders are allowed to have ties in their valuations, one can prove that monotonicity implies the existence of thresholds, while the other direction is not true in general but only under some assumption on $A_i(b_i, \mathbf{b}_{-i})$.

3.2 Characterization for Unknown Bidders

The following property generalizes the property of monotonicity of algorithms defined by [20] and characterizes truthful auctions without money and with verification.

Definition 2 *An algorithm A is k -set monotone if the following holds for any i , any \mathbf{b}_{-i} and any $a = (z, T) \in D_i$: if $A_i(a, \mathbf{b}_{-i}) = T$ then for all $b = (w, U)$ 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 see how this notion generalizes [20], it is important to understand what is U . In detail, $\sigma(T|b) = U$, in the above definition, 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, as from its definition in Sect. 2, $\sigma(T|b)$ 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.)

Theorem 2 *An algorithm A is truthful without money and with verification 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 algorithm A . Take any edge of the graph ($b = (w, \mathcal{U})$, $a = (z, \mathcal{T})$) 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 then we also have that $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$ and then 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: this means that whenever $w(U) \geq z(T)$ then it must be $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$ and therefore A is k -set monotone. The theorem follows from Proposition 1. \square

Observe, that our characterization of Theorem 2 for unknown single-minded bidders implies the existence of a threshold for any set. Namely, let A be a given 1-set monotone algorithm, and let i be a fixed bidder with declaration $(z, \mathcal{T}) \in D_i$. Then for the set $T \in \mathcal{T}$ (here, $|\mathcal{T}| = 1$), algorithm A is monotone with respect to $z(T)$ and thus there exists a critical threshold. It is not hard to see that thresholds exist also for unknown k -minded bidders, with $k > 1$.

The result in Theorem 2 also relates to the characterization of truthful CAs with money and no verification (see, e.g., Proposition 9.27 in [26]). 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 Our Characterizations

We discuss here two conceptually relevant consequences of our results above. In a nutshell, a reasonably large class of truthful mechanisms with money can be turned into truthful mechanisms without money, by using the verification paradigm.

3.3.1 Single-Minded Versus Multi-minded Bidders

Observe that our characterization of truthful mechanisms without money for CAs with 1-minded bidders with known and unknown bidders is exactly the same as the characterization of truthful mechanisms with money in this setting, see, e.g., pages 274–275 in [26]. This means that the two classes of truthful mechanisms in fact coincide. More formally, we have:

Proposition 2 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 CAs

It is known that a slight generalization of monotonicity of [20] is a sufficient property to obtain truthful mechanisms with money also for problems involving *generalized single-minded bidders* [3]. Intuitively, generalized single-minded bidders have k private numbers associated to 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 [3] for details. By Theorem 2, all the truthful mechanisms with money designed for this quite general type of bidders can be turned into truthful mechanisms without money, when the verification paradigm is justifiable. As a direct corollary of our characterization, we then have a host of truthful mechanisms without money and with verification for the (multi-objective optimization) problems studied in [3, 12].

4 Upper Bounds for Unknown Bidders

In this section we present our upper bounds for CAs with unknown k -minded bidders.

4.1 CAs with Arbitrary Supply of Goods

In this section, we consider the more general case in which elements in \mathbf{U} are available in b copies each. Note that the characterizations above hold also in this multi-unit case (i.e., it is, in fact, harmless to substitute sets of goods with multi-sets and adjust the notations accordingly so that, for example, $A_i(\mathbf{b})$ is the multi-set of goods that auction A grants to i on input \mathbf{b}). 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 [19] to our setting without money. The algorithm considers bidders in an arbitrary

Algorithm 1: Multiplicative price update algorithm

- 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 Set $S_i := \operatorname{argmax} \{v_i(S) : S \in \mathcal{S}_i \text{ such that } v_i(S) \geq \sum_{e \in S} p_e^i\}$.
 - 4 Update for each good $e \in S_i$: $p_e^{i+1} := p_e^i \cdot r$.
 - 5 Return $S = (S_1, S_2, \dots, S_n)$.
-

given order. We assume that the algorithm is given a parameter $\mu \geq 1$ such that $\mu/2 \leq v_{\max} < \mu$. We will assume that such μ is known to the mechanism, and afterwards we will modify our mechanism and show 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. For a detailed description see Algorithm 1.

Algorithm 1 is very similar to the basic algorithm of [19] (and the same happens for the analysis). The only important difference is that in [19], since they have CAs with money, the prices p_e^i are used to determine the payment for each bundle. So, in [19], each bidder i is allocated a set S_i in her demand set \mathcal{S}_i that maximizes $v_i(S_i) - \sum_{e \in S_i} p_e^i$, i.e., bidder’s i valuation for S_i minus her payment for S_i . In our case, since we do not have monetary transfers, but we use verification instead, the prices p_e^i are used to determine the threshold of each bundle. Therefore, each bidder i is now allocated a set S_i in \mathcal{S}_i that maximizes $v_i(S_i)$ among all sets $S' \in \mathcal{S}_i$ with $v_i(S') \geq \sum_{e \in S'} p_e^i$. This simple (but essential) change makes the basic algorithm of [19] truthful without money and with verification and allows us to apply the main steps of their analysis.

Let ℓ_e^i be the number of copies of good $e \in U$ allocated to all bidders preceding bidder i and $\ell_e^* = \ell_e^{n+1}$ denote the total allocation of good e to all bidders. Let, moreover, $p_e^* = p_0 \cdot r^{\ell_e^*}$ be good e ’s price at the end of the algorithm.

We claim now that if p_0 and r are chosen so that $p_0 r^b = \mu$, then the allocation $S = (S_1, \dots, S_n)$ output by Algorithm 1 is feasible, that is, it assigns at most b copies of each good to the bidders. The argument is as follows. Consider any good $e \in U$. Notice that when the b -th copy of good e is sold to any bidder then its price is updated to $p_0 r^b = \mu > v_{\max}$. Thus, good e alone has a price which is above the maximum valuation of any bidder, and so no further copy will be sold.

Next we prove two lower bounds on the social welfare $v(S) = \sum_{i=1}^n v_i(S_i)$ of the sets S_1, \dots, S_n chosen by Algorithm 1. Let OPT denote the optimal social welfare, and recall that p_e^* denotes the final price of good $e \in U$.

Lemma 2 *It holds $v(S) \geq \frac{1}{r-1} (\sum_{e \in U} p_e^* - m p_0)$ and $v(S) \geq \text{OPT} - b \sum_{e \in U} p_e^*$.*

Proof We first prove the first bound. Because the algorithm sells only sets above their thresholds, $v_i(S_i) \geq \sum_{e \in S_i} p_e^i$. Hence,

$$v(S) \geq \sum_{i=1}^n \sum_{e \in S_i} p_e^i = \sum_{i=1}^n \sum_{e \in S_i} p_0 r^{\ell_e^i} = p_0 \sum_{e \in U} \sum_{k=0}^{\ell_e^* - 1} r^k = p_0 \sum_{e \in U} \frac{r^{\ell_e^*} - 1}{r - 1},$$

which gives the bound.

For the second bound of the lemma, consider an optimal feasible allocation $T = (T_1, \dots, T_n)$. The set choice in line 3 in the algorithm of [19] is done by asking the

demand oracle.⁶ We observe here that the set choice in this line in Algorithm 1 is sufficient for showing this claim. Namely, if $v_i(T_i) \geq \sum_{e \in T_i} p_e^i$, then by the choice of the algorithm $v_i(S_i) \geq v_i(T_i)$, and so we have $v_i(S_i) \geq v_i(T_i) - \sum_{e \in T_i} p_e^i$. If on the other hand $v_i(T_i) < \sum_{e \in T_i} p_e^i$, then $v_i(S_i) \geq v_i(T_i) - \sum_{e \in T_i} p_e^i$ also holds.

Because $p_e^* \geq p_e^i$, for every i and e , this implies $v_i(S_i) \geq v_i(T_i) - \sum_{e \in T_i} p_e^*$. Summing over all bidders gives $v(S) = \sum_{i=1}^n v_i(S_i) \geq \sum_{i=1}^n v_i(T_i) - \sum_{i=1}^n \sum_{e \in T_i} p_e^* \geq v(T) - b \sum_{e \in U} p_e^*$, where the latter equation uses the fact that T is feasible so that each good is given to at most b sets. □

Combining the above lemmas yields the following result for the algorithm.

Theorem 3 *Algorithm 1 with $p_0 = \frac{\mu}{4bm}$ and $r = (4bm)^{1/b}$ produces a feasible allocation S such that $v(S) \geq \frac{\text{OPT}}{2(b(r-1)+1)} \geq \frac{\text{OPT}}{O(b \cdot (m)^{1/b})}$.*

Proof Feasibility follows from the fact that $p_0 r^b = \mu$. The first bound of Lemma 2 gives $b(r-1)v(S) \geq b \sum_{e \in U} p_e^* - bmp_0$, which by the second bound is $b(r-1)v(S) \geq b \sum_{e \in U} p_e^* - bmp_0 \geq \text{OPT} - v(S) - bmp_0 \geq \text{OPT}/2 - v(S)$, where the last inequality follows by $v_{\max} \leq \text{OPT}$. This finally gives us $v(S) \geq \frac{\text{OPT}}{2(b(r-1)+1)}$, implying the final approximation ratio. □

Theorem 4 *Algorithm 1 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 2, take two declarations of bidder i , $a = (z, T)$ and $b = (w, U)$ with $w(U) \geq z(T)$, where $T = A_i(a, \mathbf{b}_{-i})$ and $U = \sigma(T|b)$. (In this proof, A denotes Algorithm 1.) Recall that $U \subseteq T$ and $U \in \mathcal{U}$.

Note that the ordering is independent of the bids and then when i is considered the prices p_e^i for the elements e of U are the same in both $A(a, \mathbf{b}_{-i})$ and $A(b, \mathbf{b}_{-i})$. Since $T = A_i(a, \mathbf{b}_{-i})$, we note that $z(T) \geq \sum_{e \in T} p_e^i$. This yields, $w(U) \geq z(T) \geq \sum_{e \in T} p_e^i \geq \sum_{e \in U} p_e^i$. This implies that when $A(b, \mathbf{b}_{-i})$ executes line 3, the set U is taken into consideration and we can therefore conclude that $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$. This shows that A is k -set monotone and then, by Theorem 2, our claim follows. □

We now modify Algorithm 1 in order to remove the assumption on the knowledge of μ . The idea behind the modified algorithm is that we elicit the information about the maximum valuation of any bidder, which is needed for the computation of μ , by ranking first the bidder j with the maximum valuation v_{\max}^j and offering him the corresponding bundle. Since bidder j cannot overbid on this bundle, we elicit v_{\max}^j in a truthful way. The modified algorithm is presented as Algorithm 2. We have the following result.

Theorem 5 *Algorithm 2 is a truthful mechanism without money and with verification for CAs with unknown k -minded bidders. Its approximation ratio is $O(b \cdot (m)^{1/b})$.*

⁶ Given goods' prices, the demand oracle of a given bidder with valuation function v_i outputs the set S maximizing the difference between the set's valuation $v_i(S)$ and the sum of the prices of its goods.

Algorithm 2: Modified multiplicative price update algorithm

- 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 $e \in U$ do $p_e^j := p_0$.
- 5 Let for any $i = j, 1, 2, 3, 4, \dots, j - 1, j + 1, \dots, n$, $\text{next}(i)$ 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 Set $S_i := \operatorname{argmax} \{v_i(S) : S \in \mathcal{S}_i \text{ such that } v_i(S) \geq \sum_{e \in S} p_e^i\}$.
 - 8 Update for each good $e \in S_i$: $p_e^{\text{next}(i)} := p_e^i \cdot r$.
- 9 Return $S = (S_1, S_2, \dots, S_{j-1}, S_j, S_{j+1}, \dots, S_n)$.

Proof Approximation ratio and feasibility of the produced solution follow from the choice of p_0 and from setting $\mu = (1 + \epsilon)v_{\max}^j$, for $0 < \epsilon \ll 1$. Indeed, we can use the previous analysis of Algorithm 1 that did not make any assumption on the order in which bidders are processed and only required $\mu/2 \leq v_{\max} < \mu$.

We will argue now about truthfulness of the modified algorithm. Let us call bidder j in Algorithm 2, the *max bidder*. We first observe that bidder j is allocated the set in her (reported) demand with highest (reported) valuation. This is because her declaration of v_{\max}^j for her best set, say Q , is larger than $\sum_{e \in Q} p_e^j = |Q| \cdot p_0 = |Q| \cdot \frac{(1+\epsilon)v_{\max}^j}{4bm}$ since $|Q| \leq m$. Now, fix i and \mathbf{b}_{-i} . As in Definition 2, take two declarations of bidder i , $a = (z, T)$ and $b = (w, U)$ with $w(U) \geq z(T)$, where $T = A_i(a, \mathbf{b}_{-i})$ and $U = \sigma(T|b)$. (In this proof, A denotes Algorithm 2.) Recall that $U \subseteq T$ and $U \in \mathcal{U}$. Let j_a (resp., j_b) be the max bidder for the bid vector (a, \mathbf{b}_{-i}) (resp., (b, \mathbf{b}_{-i})). We distinguish three cases.

Case 1: $i = j_a$. In this case, $z(T)$ is larger than all the valuations in \mathbf{b}_{-i} . Since $w(U) \geq z(T)$ and since \mathbf{b}_{-i} is unchanged then $w(U)$ is also larger than all the valuations in \mathbf{b}_{-i} , which yields $i = j_b$. But then, as observed above, since i is the max bidder in (b, \mathbf{b}_{-i}) she will get her best set in U and therefore $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$.

Case 2: $i = j_b$. Since i is the max bidder in (b, \mathbf{b}_{-i}) then we can argue, as above, that she will get her best set and so we have $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$.

Case 3: $i \neq j_a, j_b$. Since the other bids are unchanged, in this case, we have $j_a = j_b$. This implies that the ordering in which bidders are considered is the same in both $A(a, \mathbf{b}_{-i})$ and $A(b, \mathbf{b}_{-i})$ which in turns implies that the prices p_e^i considered by the algorithm in line 7 are the same in both instances. We can then use the same arguments used in the proof of Theorem 4 to conclude that $w(A_i(b, \mathbf{b}_{-i})) \geq w(U)$.

In all the three cases we have shown that the algorithm is k -set monotone and then the claim follows from Theorem 2. □

4.1.2 Randomized Truthful CAs

We show here how to use Algorithm 2 to obtain randomized universally truthful mechanisms without money and with verification for CAs with unknown k -minded bidders with expected approximation ratios of $O(d^{1/b} \log(bm))$ and $O(m^{1/(b+1)} \log(bm))$, respectively.

Algorithm 3: Multiplicative price update algorithm with randomized rounding.

- 1 For each good $e \in \mathbf{U}$ do $p_e^1 := p_0, b_e^1 := b$.
- 2 For each bidder $i = 1, 2, \dots, n$ do
- 3 Let $\mathbf{U}_i = \{e \in \mathbf{U} \mid b_e^i > 0\}$.
- 4 Set $S_i := \operatorname{argmax} \{v_i(S) : S \in \mathcal{S}_i \text{ such that } S \subseteq \mathbf{U}_i \text{ and } v_i(S) \geq \sum_{e \in S} p_e^i\}$.
- 5 Update for each good $e \in S_i: p_e^{i+1} := p_e^i \cdot r$.
- 6 With probability q set $R_i := S_i$ else $R_i := \emptyset$.
- 7 Update for each good $e \in R_i: b_e^{i+1} := b_e^i - 1$.
- 8 Return $R = (R_1, R_2, \dots, R_n)$.

Algorithm 4: Modified multiplicative price update algorithm with randomized rounding

- 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 $p_e^j := p_0, b_e^j := b$.
- 5 Let for any $i = j, 1, 2, 3, 4, \dots, j-1, j+1, \dots, n$, $\operatorname{next}(i)$ be the next number in this order, e.g., $\operatorname{next}(j) = 1, \operatorname{next}(1) = 2, \dots, \operatorname{next}(j-1) = j+1, \dots, \operatorname{next}(n-1) = n, \operatorname{next}(n) = n+1$.
- 6 For each bidder $i = j, 1, 2, \dots, j-1, j+1, \dots, n$ do
- 7 Let $\mathbf{U}_i = \{e \in \mathbf{U} \mid b_e^i > 0\}$.
- 8 Set $S_i := \operatorname{argmax} \{v_i(S) : S \in \mathcal{S}_i \text{ such that } S \subseteq \mathbf{U}_i \text{ and } v_i(S) \geq \sum_{e \in S} p_e^i\}$.
- 9 Update for each good $e \in S_i: p_e^{\operatorname{next}(i)} := p_e^i \cdot r$.
- 10 If $i = j$ then set $R_i := S_i$ else (with probability q set $R_i := S_i$ else $R_i := \emptyset$).
- 11 Update for each good $e \in R_i: b_e^{\operatorname{next}(i)} := b_e^i - 1$.
- 12 Return $R = (R_1, R_2, \dots, R_{j-1}, R_j, R_{j+1}, \dots, R_n)$.

Observe first that if we execute Algorithm 1 with a smaller update factor $r = 2^{1/b}$, as in [19], then the output solution is infeasible and as proved in [19, Lemma 1], it allocates at most sb copies of each good to the bidders, where $s = \log(4bm)$. This simply follows from the fact that if sb copies of good $e \in \mathbf{U}$ were sold, then its price is $p_0 2^{\log(4bm)} = \mu > v_{\max}$. 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 3 indeed implies an $O(1)$ -approximation (see also [19, Theorem 1]). This idea leads to the following randomized algorithm in [19]: 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 (see also Algorithm 3 for a precise description). They prove that this algorithm is a universally truthful $O(d^{1/b} \log(4bm))$ -approximation mechanism.

We now introduce the same randomization idea into our Algorithm 2. The resulting algorithm is Algorithm 4, where we assume $r = 2^{1/b}$.

Theorem 6 *Algorithm 4 is a universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders. Its expected approximation ratio is $O(d^{1/b} \cdot \log(bm))$.*

Proof Approximation guarantee and feasibility of the output solution R follows from essentially the same arguments and in the same way as in [19]. We will argue now about universal truthfulness of Algorithm 4. This algorithm can be viewed as a probability distribution over deterministic algorithms. Each such algorithm, call it A , is defined by a 0/1-vector $a \in \{0, 1\}^{n-1}$ and first selects and serves the max bidder j and then serves the remaining $n - 1$ bidders $1, 2, \dots, j - 1, j + 1, \dots, n$. When serving bidder $i \neq j$, algorithm A deterministically allocates set S_i to bidder $i < j$ if and only if $a_i = 1$ and to bidder $i > j$ if and only if $a_{i-1} = 1$. Consider algorithm A where vector $a = (1, 1, \dots, 1)$ and observe that this corresponds to our Algorithm 2. Thus, to show that A is (deterministically) truthful we use the same argument of the proof of Theorem 5 and the additional observation that bidders whose corresponding bit in the vector a is 0 have no incentive to lie, since they are not served anyway. \square

Finally we can also obtain a universally truthful mechanism in case demanded sets have unbounded sizes.

Theorem 7 *There exist a universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders with an expected approximation ratio of $O(m^{1/(b+1)} \cdot \log(bm))$.*

Proof We use Algorithm 4 as a subroutine and use a standard randomization on the top of this algorithm. Namely, the mechanism flips a fair coin to choose one out of two algorithms. If the coin shows head, then Algorithm 4 is executed with parameter q set analogously to the case of sets of size at most $d = \lfloor m^{\frac{b}{b+1}} \rfloor$, that is, with $q^{-1} = 2ed^{1/b} \log(4bm)$. If the coin shows tail, then the mechanism only considers sets of full cardinality m . This setting corresponds to an auction where each bidder wants to buy only a single super item (corresponding to U) which is available in b copies. The b copies of the super item are sold by calling Algorithm 4 with $m = 1$ (single item) and $d = 1$ (sets of cardinality 1). This mechanism can easily be shown to be universally truthful and a simple analysis shows the claimed approximation ratio, see, e.g., [19]. \square

4.2 CAs with Single Supply

We now go back to the case in which the goods in U are provided with single supply. 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 now present a simple greedy algorithm for CAs where the supply $b = 1$, see Algorithm 5. (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 [17].) Recall that each bidder $i = 1, 2, \dots, n$ declares (v_i, S_i) , where S_i is a collection of k sets bidder i demands and $v_i(S)$ is the valuation of set $S \in S_i$. The algorithms first organizes

Algorithm 5: 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 $\mathcal{P} := \mathcal{P} \cup \{S_i\}$.
- 8 $\mathcal{B} := \mathcal{B} \cup \beta(i)$.
- 9 Return \mathcal{P} .

all these pairs in non-decreasing order of valuations (using bidders' ids to break ties, if any). Observe that sets S_1, \dots, S_l are all the sets demanded by all bidders (with non-zero bids), i.e., $\{S_1, \dots, S_l\} = S_1 \cup \dots \cup S_n$. Subsequently, it scans this list and greedily assigns sets to bidders while maintaining the feasibility of the allocation.

We will use the linear programming duality theory to prove the approximation guarantees of our algorithm. Let us denote the set family $\mathcal{S} = \cup_{i=1}^n S_i$, where bidder i demands sets 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 $[n]$ be the set $\{1, \dots, n\}$. The standard LP relaxation of our problem is (see e.g., [26, Section 11.3.1]):

$$\max \sum_{i=1}^n \sum_{S \in \mathcal{S}_i} b_i(S) x_i(S) \tag{6}$$

$$\text{s.t. } \sum_{i=1}^n \sum_{S: S \in \mathcal{S}_i, e \in S} x_i(S) \leq 1 \quad \forall e \in \mathbf{U} \tag{7}$$

$$\sum_{S \in \mathcal{S}_i} x_i(S) \leq 1 \quad \forall i \in [n] \tag{8}$$

$$x_i(S) \geq 0 \quad \forall i \in [n] \forall S \in \mathcal{S}_i, \tag{9}$$

The corresponding dual linear program is then the following:

$$\min \sum_{e \in \mathbf{U}} y_e + \sum_{i=1}^n z_i \tag{10}$$

$$\text{s.t. } z_i + \sum_{e \in S} y_e \geq b_i(S) \quad \forall i \in [n] \quad \forall S \in \mathcal{S}_i \tag{11}$$

$$z_i, y_e \geq 0 \quad \forall i \in [n] \quad \forall e \in \mathbf{U}. \tag{12}$$

In this dual linear program dual variable z_i corresponds to the constraint (8).

Theorem 8 *Algorithm 5 is a $\min\{m, d + 1\}$ -approximation algorithm for CAs with k -minded bidders.*

Proof Suppose that Algorithm 5 has terminated and output solution \mathcal{P} . Let $\mathbf{U}(\mathcal{P}) = \cup_{S \in \mathcal{P}} S$. Notice that for each set $S \in \mathcal{S}$ that was not chosen to the final solution \mathcal{P} , there either is an element $e \in \mathbf{U}(\mathcal{P}) \cap S$ which was the *witness* of that event during the execution of the algorithm, 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 $\mathbf{U}(\mathcal{P})$ one witness for S . In case if there is more than one witness in $\mathbf{U}(\mathcal{P}) \cap S$, we keep in $\mathbf{U}(\mathcal{P})$ the (arbitrary) witness for S

that belongs to the set among sets $\{T \in \mathcal{P} : \mathbf{U}(\mathcal{P}) \cap S \cap T \neq \emptyset\}$ that was considered first by the greedy order. We discard the remaining elements from $\mathbf{U}(\mathcal{P})$.

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

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.

We then assume that $m \geq 2$. We now define a dual solution during the execution of Algorithm 5. We need to know the output solution \mathcal{P} for the definition of this dual solution, which is needed only for analysis. In line 4 of Algorithm 5 we initialize these variables: $y_e := 0$ for all $e \in \mathbf{U}$ and $z_i := 0$ for all $i \in [n]$. We add the following in line 7 of Algorithm 5: $y_e := \Delta_e^{S_i}$, 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 $e \in S_i \setminus \mathbf{U}(\mathcal{P})$ the value of y_e is not updated and remains zero. We also add the following instruction in line 7 of Algorithm 5: $z_{\beta(i)} := b_{\beta(i)}(S_i)$.

It is obvious that the dual solution provides a lower bound on the cost of the output solution:

$$\sum_{e \in \mathbf{U}} y_e \leq \sum_{S_i \in \mathcal{P}} b_{\beta(i)}(S_i). \tag{13}$$

We will show now 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 constraints (11) are fulfilled for all sets $S \in \mathcal{S}$. Thus, we have to prove that, for each set $S \in S \cap S_i$,

$$z_i + d' \sum_{e \in S} y_e \geq b_i(S). \tag{14}$$

Suppose first that $S = S_r \in \mathcal{S} \setminus \mathcal{P}$, and let $\beta(r) = i$. There are two possible reasons that set S has not been included in the solution \mathcal{P} : (i) **Case (a)**: there must be an element $e \in \mathbf{U}(\mathcal{P})$ such that $e \in S$, or (ii) **Case (b)**: there is another set $S' \in \mathcal{P}$ with $S, S' \in S_i$.

Let us first consider Case (a). In that case adding set S to solution \mathcal{P} would violate constraint (7). 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 of $|S| = m$, that is, $S = \mathbf{U}$, 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 (14) in Case (a).

We consider now Case (b). Suppose that $S = S_r \in \mathcal{S} \setminus \mathcal{P}$ and there is another set $S' = S_j \in \mathcal{P}$ with $S, S' \in S_i$. In this case we have $i = \beta(j) = \beta(r)$. Observe that when set S' was chosen by Algorithm 5 the dual variable z_i was updated in line 7 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 (14) in this case.

Notice that claim (14) follows immediately from the definition of z_i if set $S \in S_i$ has been chosen by our algorithm, that is, $S \in \mathcal{P}$. This concludes the proof of (14).

Finally, we put all the pieces together. We have shown that the dual solution $(d' \cdot y, z)$ is feasible for the dual linear program and so by weak duality $\sum_{i=1}^n z_i + d' \sum_{e \in \mathbf{U}} y_e$

is an upper bound on the value of the optimal integral solution to our problem. We have also shown in (13), that $\sum_{e \in U} y_e \leq \sum_{S_i \in \mathcal{P}} b_{\beta(i)}(S_i)$. Therefore, by letting OPT denote the optimal social welfare, we obtain that

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

□

We now prove the truthfulness of Algorithm 5.

Theorem 9 *Algorithm 5 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 2, take two declarations of bidder i , $a = (z, T)$ and $b = (w, U)$ with $w(U) \geq z(T)$, where $T = A_i(a, \mathbf{b}_{-i})$ and $U = \sigma(T|b)$. (In this proof, A denotes Algorithm 5.) Recall that $U \in \mathcal{U}$ and $U \subseteq T$.

Let \mathbf{S}_a (respectively, \mathbf{S}_b) be the set comprised of the sets in declarations of \mathbf{b}_{-i} processed by $A(a, \mathbf{b}_{-i})$ (respectively, $A(b, \mathbf{b}_{-i})$) when $z(T)$ (respectively, $w(U)$) is considered. Since A grants T to bidder i in the instance (a, \mathbf{b}_{-i}) then it must be the case that $T \cap S = \emptyset$ for all $S \in \mathbf{S}_a$ granted by A . Since $w(U) \geq z(T)$, then we have that $\mathbf{S}_b \subseteq \mathbf{S}_a$. Thus, since $U \subseteq T$ then we have that $U \cap S = \emptyset$ for all $S \in \mathbf{S}_b$ 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 2.

Borodin and Lucier [1, Theorem 2] proved a lower bound of $\Omega(m)$ on the approximation ratio of any truthful greedy priority mechanism with money for instances with demanded sets of cardinality at most 2. Nevertheless, we have shown that Algorithm 5 is truthful and achieves an approximation ratio of at most 3 for such instances. We here investigate the reasons behind this sharp contrast.

Proposition 3 *There are no payments that augment Algorithm 5 so to make a truthful mechanism for CAs with k -minded bidders, even in the case of known bidders.*

Proof We consider 3 instances I, I', I'' , with known bidders and $\mathbf{U} = \{a, b\}$. In all of them, bidder 1 is interested in $\mathcal{S}_1 = \{\{a\}, \{b\}\}$, and bidder 2 is interested in $\mathcal{S}_2 = \{\{a\}\}$ and has $v_2(\{a\}) = 1$. Let $\delta > 0$ be some small constant less than $1/2$. In I , bidder 1 has $v_1(\{a\}) = 1 - \delta$ and $v_1(\{b\}) = 0$. In I' , bidder 1 has $v'_1(\{a\}) = 1 + \delta$ and $v'_1(\{b\}) = 1$. In I'' , bidder 1 has $v''_1(\{a\}) = 1 - \delta$ and $v''_1(\{b\}) = 1 - 2\delta$. Therefore, the allocation of Algorithm 5 is $\{b\}$ to bidder 1 and $\{a\}$ to bidder 2 for instance I , $\{a\}$ to bidder 1 and \emptyset to bidder 2 for instance I' , and $\{b\}$ to bidder 1 and $\{a\}$ to bidder 2 for instance I'' . Now we consider the vertices corresponding to v_1, v'_1 , and v''_1 of the declaration graph for Algorithm 5 and find an edge (v_1, v'_1) of weight $-(1 - \delta)$, an edge (v'_1, v''_1) of weight δ , and an edge (v''_1, v_1) of weight 0. Therefore, the declaration graph contains a negative cycle $v_1 \rightarrow v'_1 \rightarrow v''_1 \rightarrow v_1$, which implies that the valuation-greedy mechanism cannot be truthfully implemented with money. □

4.2.2 Randomized Exponential-Time Mechanism

We describe the exponential-time mechanism, or RandExp in brief. Let I be an instance of CAs with unknown k -minded bidders, and let I_ℓ , $1 \leq \ell \leq k$, be the subinstance of I that consists of the elementary bids $(i, S_i^\ell, v_i(S_i^\ell))$, $i \in N$, where S_i^ℓ denotes 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 10 *RandExp is a universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders. It achieves an approximation ratio of k .*

Proof Since all bidders are single-minded in each subinstance I_ℓ , the allocation corresponding to the maximum social welfare OPT_ℓ using a fixed tie-breaking rule is, by Theorem 2, truthful, for each $\ell \in [k]$. Therefore, RandExp is universally truthful. As for the approximation ratio, the expected social welfare of RandExp is $\sum_{\ell=1}^k \text{OPT}_\ell / k$. Since the maximum social welfare of I is at most $\sum_{\ell=1}^k \text{OPT}_\ell$, RandExp has an approximation ratio of k . \square

4.2.3 Randomized Polynomial-Time Mechanism

We describe the polynomial-time 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} a set with valuation v_{\max} . Moreover, let I_s be the subinstance of I that consists of the elementary bids $(i, S, v_i(S))$, $i \in N$, where for each set S , $|S| \leq \sqrt{m}$. 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 5 on the subinstance I_s . Next, we show the following fact.

Theorem 11 *RandPoly is a universally truthful mechanism without money and with verification for CAs with unknown k -minded bidders. It achieves an approximation ratio of $O(\sqrt{m})$.*

Proof As for the truthfulness of RandPoly, the deterministic allocation of the set of maximum valuation with a fixed tie-breaking rule is, by Theorem 2, truthful. Furthermore, the formation of the subinstance I_s and the application of Algorithm 5 to it can be regarded as a run of the algorithm on I where any elementary bid $(i, S, v_i(S))$ with $|S| > \sqrt{m}$ is considered infeasible and immediately rejected. By using the arguments used in Theorem 9, we can prove that this \sqrt{m} -cardinality-sensitive variant of Algorithm 5 to I is truthful. Therefore, RandPoly is universally truthful.

As for the approximation ratio, we let OPT denote the optimal social welfare of I , OPT_s denote the optimal social welfare of the subinstance I_s , and OPT_l denote the optimal social welfare of the subinstance $I_l = I \setminus I_s$. We observe that $\text{OPT} \leq \text{OPT}_s + \text{OPT}_l$. Since I_l contains only elementary bids $(i, S, v_i(S))$ with $|S| > \sqrt{m}$ and

$v_i(S) \leq v_{\max}$, $\text{OPT}_I \leq \sqrt{m}v_{\max}$. Moreover, by Theorem 8, the allocation computed by the application of valuation-greedy to I_s is a $(\sqrt{m} + 1)$ -approximation of OPT_s . Hence, the expected social welfare of RandPoly is at least:

$$\frac{1}{2} \left(v_{\max} + \frac{\text{OPT}_s}{\sqrt{m} + 1} \right) \geq \frac{\text{OPT}_s + \text{OPT}_I}{2(\sqrt{m} + 1)}$$

Therefore, the approximation ratio of RandPoly is $2(\sqrt{m} + 1)$. \square

5 Lower Bounds for Known Bidders

In this section, we prove lower bounds on the approximation ratio of (deterministic or randomized) mechanisms for CAs with known k -minded bidders. The main message of our results is that unlike the setting of CAs with money, where if we have exponential computational time, we can truthfully implement the optimal solution using VCG payments, in our setting with known k -minded bidders, for $k \geq 2$, it is not possible to truthfully implement the optimal solution, even in exponential time.

5.1 Deterministic Mechanisms

We first adapt the proof of [8, 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 does not make any assumptions on the computational power of the mechanism, and holds even for exponential-time mechanisms.

Theorem 12 *There are no deterministic truthful mechanisms with approximation ratio better than 2 for CAs with known 2-minded bidders. This holds even for simple instances with $n = 2$ bidders, and $m = 2$ goods.*

Proof For sake of contradiction, let us assume that there is a deterministic truthful mechanism A with an approximation ratio of $2 - \delta$, for some $\delta > 0$. We consider two instances where $\mathbf{U} = \{a, b\}$, and both bidders are in $\mathcal{S}_1 = \mathcal{S}_2 = \{\{a\}, \{b\}\}$. In the first instance, $v_1(\{a\}) = 1 + \delta$ and $v_1(\{b\}) = 1$, and $v_2(\{a\}) = 1 + \delta$ and $v_2(\{b\}) = 1$. Since A is a $(2 - \delta)$ -approximation algorithm, $A(v_1, v_2)$ must allocate both sets $\{a\}$ and $\{b\}$. Without loss of generality, we assume that $A(v_1, v_2)$ allocates $\{a\}$ to bidder 1 and $\{b\}$ to bidder 2. Moreover, by Theorem 1, if $v'_2(\{a\}) = 1 + \delta$ and $v'_2(\{b\}) = 0$, then $A(v_1, v'_2)$ must allocate $\{a\}$ to bidder 1 and $\{b\}$ to bidder 2. Therefore, A has an approximation ratio of at least $(2 + \delta)/(1 + \delta)$, which is larger than $2 - \delta$, for all $\delta > 0$. \square

We note here, that our Algorithm 5 is a truthful 2-approximate mechanism on instances used in the proof of Theorem 12. This theorem 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 consideration. Furthermore, it shows that, unlike the case of single-minded bidders, 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.

5.2 Universally Truthful Mechanisms

Next, we apply Yao’s principle and show that no randomized mechanism that is universally truthful can achieve an approximation ratio better than $5/4$.

Theorem 13 *There are no randomized mechanisms that are universally truthful and have approximation ratio better than $5/4$ for CAs with known 2-minded bidders. This holds even for simple instances with $n = 2$ bidders, and $m = 2$ goods.*

Proof We present a probability distribution over instances of CAs with $n = 2$ known 2-minded bidders, and $m = 2$ goods, for which the best deterministic truthful mechanism has expected approximation ratio greater than $5/4 - \delta$, for any $\delta > 0$. We consider two instances I and I' where $\mathbf{U} = \{a, b\}$, the first bidder is interested in $\mathcal{S}_1 = \{\{a, b\}, \{b\}\}$, and the second bidder 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 A , applied to instance I , allocate $\{a, b\}$ to bidder 1 and \emptyset to bidder 2. Then, by Theorem 1, 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 instance I , does not allocate $\{a, b\}$ to bidder 1, its expected social welfare is at most $(4 - \delta)/2$, and its expected approximation ratio is $(5 - \delta)/(4 - \delta) > 5/4 - \delta$, a contradiction. \square

5.3 Truthful in Expectation Mechanisms

Finally, we show a weaker lower bound of 1.09 on the approximation ratio achievable by the larger class of randomized mechanisms that are truthful in expectation.

Theorem 14 *There are no randomized mechanisms that are truthful in expectation and have approximation ratio less or equal than 1.09 for CAs with known 2-minded bidders. This holds even for simple instances with $n = 2$ bidders, and $m = 2$ goods.*

Proof For sake of contradiction, we assume that there is a randomized truthful-in-expectation mechanism A with approximation ratio at most $\rho = 1.09$.

As in the proofs of Theorems 12 and 13, we consider instances where $\mathbf{U} = \{a, b\}$, the first bidder is interested in $\mathcal{S}_1 = \{\{a, b\}, \{b\}\}$, and the second bidder is interested in $\mathcal{S}_2 = \{\{a\}\}$. We consider two instances I and I' . In both of them, the valuation of bidder 2 is $v_2(\{a\}) = 1$. The valuation of bidder 1 is $v_1(\{a, b\}) = \varphi$, where $\varphi = (1 + \sqrt{5})/2$ is the golden ratio, and $v_1(\{b\}) = 0$ in I , and $v'_1(\{a, b\}) = \varphi$ and $v'_1(\{b\}) = 1$ in I' .

We assume that A can have only two different solutions for instances I and I' . More specifically, either A allocates $\{a, b\}$ to bidder 1 and \emptyset to bidder 2, which happens with probability p for instance I and q for instance I' , or A allocates $\{b\}$ to bidder 1 and $\{a\}$ to bidder 2, which happens with probability $1 - p$ for instance I and $1 - q$ for instance I' . Note that this assumption is without loss of generality, since all the other feasible solutions have worst social welfare than the two considered.

Using that the approximation ratio of A is at most ρ , we obtain that:

$$\frac{\varphi}{p\varphi + 1 - p} \leq \rho \Rightarrow p \geq \frac{\varphi - \rho}{\rho(\varphi - 1)} \quad (15)$$

$$\frac{2}{q\varphi + 2(1 - q)} \leq \rho \Rightarrow 1 - q \geq \frac{2 - \rho\varphi}{\rho(2 - \varphi)} \quad (16)$$

where (15) follows from the approximation ratio of A for instance I , and (16) follows from the approximation ratio of A for instance I' .

Moreover, since A is truthful in expectation, the expected welfare of bidder 1 from A 's allocation for instance I , which is $p\varphi$, does not exceed her expected welfare from A 's allocation for instance I' , which is $q\varphi + (1 - q)$. Otherwise, bidder 1 could underbid on $\{b\}$ by declaring v_1 , and get an expected welfare of $p\varphi$. Therefore, we obtain that:

$$p\varphi \leq q\varphi + 1 - q \Rightarrow q \geq \frac{p\varphi - 1}{\varphi - 1} \quad (17)$$

Combining (15), (16), and (17), we conclude that ρ satisfies that:

$$\frac{2 - \rho\varphi}{\rho(2 - \varphi)} + \frac{\varphi \frac{\varphi - \rho}{\rho(\varphi - 1)} - 1}{\varphi - 1} \leq 1$$

This is a contradiction, because the inequality above does not hold if $\rho \in [1, 1.09]$. Thus, we conclude that any randomized mechanism A for CAs that is truthful-in-expectation, has an approximation ratio worse than 1.09. \square

5.4 Priority Mechanisms

Our lower bounds in this section apply to truthful mechanisms that use algorithms that operate according to the priority framework introduced in [2].

We now briefly introduce this framework. The input of a priority mechanism is a finite subset I of the class \mathcal{I} of all permissible input items. For CAs, we consider two classes of input items. The first class consists of elementary bids, i.e., each item is a triple $(i, S, v_i(S))$, where i is the bidder, S is one of i 's demanded sets, and $v_i(S)$ is i 's valuation for S . The second class of input items consists of bidders, i.e., each item is a pair (i, v_i) , where i is the bidder and v_i is i 's valuation function. The output of a priority mechanism consists of a decision for each input item processed. If elementary bids are the input items, the output consists of an accept or reject decision for each bid. If a bid $(i, S, v_i(S))$ is accepted, S is allocated to i , and the algorithm obtains a welfare of $v_i(S)$. If bidders (i, v_i) are the input items, the output consists of a (possibly empty) set S allocated to bidder i , and the algorithm collects a welfare of $v_i(S)$.

A (possibly adaptive) priority mechanism A receives as input a finite set of items $I \subseteq \mathcal{I}$, and proceeds in rounds, processing a single item in each round. While there are unprocessed items in I , A selects a total order \mathcal{T} on \mathcal{I} without looking at the set

of unprocessed items. It is important that \mathcal{T} can be any total order on \mathcal{I} , and that for adaptive priority algorithms, the order may be different in each round. In each round, A receives the first (according to \mathcal{T}) unprocessed item $x \in I$ and makes an irrevocable decision for it (e.g., if x is an elementary bid, A accepts or rejects it, if x is a bidder, A decides about the set allocated to her). Then, x is removed from I .

5.4.1 Lower Bound for Priority Mechanisms Processing Elementary Bids

We first show that any truthful priority mechanism A that processes elementary bids has an approximation ratio of at least d for CAs with known bidders. The proof of the next result adapts the proof of [1, Theorem 3].

Theorem 15 *Let A be a truthful priority mechanism with verification and no money for CAs with known k -minded bidders. If A processes elementary bids then the approximation ratio of A is greater than $(1 - \delta)d$, for any $\delta > 0$.*

Proof For sake of contradiction, let us assume that for some given d , $2 \leq d \leq m$, there is a truthful priority mechanism A for CAs with known k -minded bidders that achieves an approximation ratio of $(1 - \delta)d$, for some constant $\delta > 0$.

Let L be any subset of \mathbf{U} of cardinality d . As an input to A , we consider an instance I_1 that for each bidder i , contains elementary bids $(i, L, 1 + \delta)$ and $(i, S, 1)$, for all $\emptyset \neq S \subset L$. As a priority mechanism, A selects a bid from I_1 and considers it first. In the following, we distinguish between the case where the first bid is $(i, L, 1 + \delta)$, for some bidder i , and the case where the first bid is $(i, S, 1)$, for some bidder i and some set S , and show how to arrive at contradiction in both.

Case 1. Let us assume that the first bid is $(i, L, 1 + \delta)$, for some bidder i . Then, if A accepts $(i, L, 1 + \delta)$, it obtains a social welfare of $1 + \delta$, while the optimal social welfare is d , which contradicts the hypothesis that the approximation ratio of A is $(1 - \delta)d$. If A rejects $(i, L, 1 + \delta)$, we consider A 's approximation ratio for an instance $I_2 \subseteq I_1$ that includes only the elementary bid $(i, L, 1 + \delta)$. Since A cannot distinguish between I_1 and I_2 , it rejects $(i, L, 1 + \delta)$ when considering I_2 , which leads to an unbounded approximation ratio for I_2 .

Case 2. Let us assume that the first bid is $(i, S, 1)$, for some bidder i and some set $\emptyset \neq S \subset L$. Then, if A accepts $(i, S, 1)$, we consider A 's approximation ratio for an instance $I_3 \subseteq I_1$ that includes the elementary bids $(i, S, 1)$ and $(i, L, 1 + \delta)$. Since A cannot distinguish between I_1 and I_3 , it again selects bid $(i, S, 1)$ first and accepts it, when it considers I_3 . But then consider the instance I'_3 in which i changes $(i, S, 1)$ into $(i, S, 1/d)$. Since A has an approximation ratio of $(1 - \delta)d$, it must allocate L to bidder i in I'_3 . But this contradicts the hypothesis that A is truthful since the two bids of bidder i in I_3 and I'_3 do not satisfy k -monotonicity (cf. Definition 1).

If A rejects $(i, S, 1)$, we consider A 's approximation ratio for an instance $I_4 \subseteq I_1$ that includes only the elementary bid $(i, S, 1)$. Since A cannot distinguish between I_1 and I_4 , it rejects $(i, S, 1)$ when considering I_4 , which leads to an unbounded approximation ratio for I_4 . \square

We note that with a minor change in the proof, the lower bound of Theorem 15 applies to the special case of 2-minded bidders. Thus, taking into account instances with $d = m$, we obtain the following result.

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

5.4.2 Lower Bound for Priority Mechanisms Processing Bidders

We next show that any truthful priority mechanism A that processes bidders has an approximation ratio of at least $m/2$ for CAs with known k -minded bidders. We note that such priority mechanism A is potentially more powerful than a priority mechanism that processes 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 result adapts the proof of [1, Theorem 4] to our setting.

Theorem 16 *Let A be a truthful priority mechanism with verification and no money for CAs with known 2-minded bidders. If A processes bidders then the approximation ratio of A is greater than $(1 - \delta)m/2$, for any $\delta > 0$.*

Proof For sake of contradiction, let us assume that there is a truthful priority mechanism A for CAs with known 2-minded bidders that processes bidders and achieves an approximation ratio of $(1 - \delta)m/2$, for some constant $\delta > 0$.

We consider a universe $\mathbf{U} = \{a_1, \dots, a_m\}$ and m bidders. For each i , $1 \leq i \leq m$, we let g_i be a single-minded valuation where the demanded set is $\{a_i\}$ with valuation 1. More specifically, for each $S \subseteq \mathbf{U}$, $g_i(S) = 1$, if $a_i \in S$, and $g_i(S) = 0$, otherwise. Moreover, for each i , $1 \leq i \leq m$, we let f_i be a double-minded valuation where the demanded set is either $\{a_i\}$ with valuation $m^2 + \delta$, or \mathbf{U} with valuation $m^2 + 2\delta$. More specifically, $f_i(\mathbf{U}) = m^2 + 2\delta$, and for each $S \subset \mathbf{U}$, $f_i(S) = m^2 + \delta$, if $a_i \in S$, and $f_i(S) = 0$, otherwise. In the following, we only consider *restricted instances* of CAs where every bidder has a valuation of type either g or f .

We first show that for any bidder i and for all instances where all bidders $j \neq i$ have single-minded valuations of type g and i has a valuation of type f , A allocates \mathbf{U} to i and \emptyset to any bidder $j \neq i$ (this claim is the equivalent of [1, Lemma 5] in our setting). We let f_k , for some $1 \leq k \leq m$, be the valuation of bidder i . Since the optimal social welfare is at least $m^2 + 2\delta$, A allocates either \mathbf{U} or a set $S \supseteq \{a_k\}$ to bidder i . Otherwise, the social welfare of A would be at most $m - 1$, which contradicts the hypothesis that the approximation ratio of A is $(1 - \delta)m/2$. However, the fact that A is truthful, implies, by Theorem 1, that A must assign \mathbf{U} to i on input g valuations from bidders other than i and f_k valuation from i . Indeed, assume that it is not the case and consider a new instance in which bidder i declares a single-minded valuation where the demanded set is \mathbf{U} with valuation $m^2 + 2\delta$. Because of the approximation guarantee of A , on this new instance, A must grant \mathbf{U} to i . The two declarations would then contradict k -monotonicity (cf. Definition 1). Therefore, A allocates \mathbf{U} to bidder i and \emptyset to any bidder $j \neq i$.

Using this claim, we can prove the following proposition which is identical to [1, Lemma 6]. The proof is by induction on i , and is omitted because it is essentially identical to the proof in [1, Lemma 6]. In fact, that proof of [1, Lemma 6] uses only standard properties of priority algorithms, the assumption that the approximation ratio

of A is $(1 - \delta)m/2$, and [1, Lemma 5], which, in our case, is replaced by the claim above.

Proposition 4 *Let A be any truthful priority mechanism which for restricted instances with m goods, achieves an approximation ratio of $(1 - \delta)m/2$, for some constant $\delta > 0$. Then, there exists a labeling of the bidders and the goods such that the following holds for all $i \in \{0, 1, \dots, m/2 - 1\}$. Let instance $I_i = \{(j, g_j) : 1 \leq j \leq i\}$. Then, for any restricted instance $I \supseteq I_i$, A considers all the bidders in I_i before all other bidders in I , and allocates \emptyset to each bidder in I_i .*

Proof Using Proposition 4, we can now complete the proof of the lemma. Let $I' = \{(j, g_j) : 1 \leq j \leq m/2 - 1\}$ be the instance I_i defined in Proposition 4 for $i = m/2 - 1$, and let $I = I' \cup \{(j, g_m) : m/2 \leq j \leq m\}$. We note that the optimal social welfare for I is $m/2$, and that I is a restricted instance such that $I' \subseteq I$, as required by Proposition 4. Therefore, mechanism A considers bidders $1, \dots, m/2 - 1$ first and allocates \emptyset to each of them. Since bidders $m/2, \dots, m$ are all single-minded for good a_m , the social welfare of A for instance I is at most 1, which contradicts the hypothesis that the approximation ratio of A is $(1 - \delta)m/2$. \square

5.5 Discussion

A step that seems necessary for an approximation ratio of $O(\sqrt{m})$ for CAs is that the algorithm somehow compares the social welfare and chooses the best of the following two extreme solutions: a solution that only consists of the most valuable set demanded by some bidder, and a solution consisting of many small sets with a large total valuation. Otherwise, the algorithm cannot achieve an approximation ratio of $o(m)$ even for the simple case where bidder 1 is double-minded for $U = \{a_1, \dots, a_m\}$ 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 seems 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 5, if we only allocate sets of cardinality at most \sqrt{m} , gives an $O(\sqrt{m})$ -approximation (see Theorem 11 for the analysis, see also e.g., [3, Section 6] for another example of an $O(\sqrt{m})$ -approximation algorithm based on a similar comparison).

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 seems the main obstacle towards a deterministic truthful $O(\sqrt{m})$ -approximate mechanism for CAs with k -minded bidders. So, the main open problem arising from our work is prove (or disprove) that there is a strong lower bound of $\Omega(m)$ on the approximation ratio of deterministic truthful mechanisms.

References

1. Borodin, A., Lucier, B.: On the limitations of greedy mechanism design for truthful combinatorial auctions. In: Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6–10, 2010, Proceedings, Part I, pp. 90–101 (2010)
2. Borodin, A., Nielsen, M.N., Rackoff, C.: (Incremental) priority algorithms. *Algorithmica* **37**(4), 295–326 (2003)
3. Briest, P., Krysta, P., Vöcking, B.: Approximation techniques for utilitarian mechanism design. *SIAM J. Comput.* **40**(6), 1587–1622 (2011)
4. Caragiannis, I., Elkind, E., Szegedy, M., Yu, L.: Mechanism design: from partial to probabilistic verification. In: ACM Conference on Electronic Commerce, pp. 266–283 (2012)
5. Carroll, G.: When are local incentive constraints sufficient? *Econometrica* **80**(2), 661–686 (2012)
6. Celik, G.: Mechanism design with weaker incentive compatibility constraints. *Games Econ. Behav.* **56**(1), 37–44 (2006)
7. Dobzinski, S.: Two randomized mechanisms for combinatorial auctions. In: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 10th International Workshop, APPROX 2007, and 11th International Workshop, RANDOM 2007, Princeton, NJ, USA, August 20–22, 2007, Proceedings, pp. 89–103 (2007)
8. Dughmi, S., Ghosh, A.: Truthful assignment without money. In: Proceedings 11th ACM Conference on Electronic Commerce (EC-2010), Cambridge, Massachusetts, USA, June 7–11, 2010, pp. 325–334 (2010)
9. Elkind, E., Sahai, A., Steiglitz, K.: Frugality in path auctions. In: Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11–14, 2004, pp. 701–709 (2004)
10. Fotakis, D., Tzamos, C.: Winner-imposing strategyproof mechanisms for multiple facility location games. *Theor. Comput. Sci.* **472**, 90–103 (2013)
11. Fotakis, D., Zampetakis, E.: Truthfulness flooded domains and the power of verification for mechanism design. *ACM Trans. Econ. Comput.* **3**(4), 20 (2015)
12. Grandoni, F., Krysta, P., Leonardi, S., Ventrè, C.: Utilitarian mechanism design for multiobjective optimization. *SIAM J. Comput.* **43**(4), 1263–1290 (2014)
13. Green, J.R., Laffont, J.: Partially verifiable information and mechanism design. *Rev. Econ. Stud.* **53**, 447–456 (1986)
14. Hazan, E., Safra, S., Schwartz, O.: On the complexity of approximating k-set packing. *Comput. Complex.* **15**(1), 20–39 (2006)
15. Holzman, R., Kfir-Dahav, N., Monderer, D., Tennenholtz, M.: Bundling equilibrium in combinatorial auctions. *Games Econ. Behav.* **47**, 104–123 (2004)
16. Koutsoupias, E.: Scheduling without payments. *Theory Comput. Syst.* **54**(3), 375–387 (2014)
17. Krysta, P.: Greedy approximation via duality for packing, combinatorial auctions and routing. In: Mathematical Foundations of Computer Science 2005, 30th International Symposium, MFCS 2005, Gdansk, Poland, August 29–September 2, 2005, Proceedings, pp. 615–627 (2005)
18. Krysta, P., Ventrè, C.: Combinatorial auctions with verification are tractable. *Theor. Comput. Sci.* **571**, 21–35 (2015)
19. Krysta, P., Vöcking, B.: Online mechanism design (randomized rounding on the fly). In: Automata, Languages, and Programming—39th International Colloquium, ICALP 2012, Warwick, UK, July 9–13, 2012, Proceedings, Part II, pp. 636–647 (2012)
20. Lehmann, D.J., O’Callaghan, L., Shoham, Y.: Truth revelation in approximately efficient combinatorial auctions. *J. ACM* **49**(5), 577–602 (2002)
21. Lu, P., Sun, X., Wang, Y., Zhu, Z.A.: Asymptotically optimal strategy-proof mechanisms for two-facility games. In: Proceedings 11th ACM Conference on Electronic Commerce (EC-2010), Cambridge, Massachusetts, USA, June 7–11, 2010, pp. 315–324 (2010)
22. Moore, J.: Global incentive constraints in auction design. *Econometrica* **52**(6), 1523–1535 (1984)
23. Mu’alem, A., Nisan, N.: Truthful approximation mechanisms for restricted combinatorial auctions. *Games Econ. Behav.* **64**(2), 612–631 (2008)
24. Nisan, N.: The communication complexity of approximate set packing and covering. In: Automata, Languages and Programming, 29th International Colloquium, ICALP 2002, Malaga, Spain, July 8–13, 2002, Proceedings, pp. 868–875 (2002)
25. Nisan, N., Ronen, A.: Algorithmic mechanism design. *Games Econ. Behav.* **35**, 166–196 (2001)

26. Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.: *Algorithmic Game Theory*. Cambridge University Press, New York (2007)
27. Nissim, K., Smorodinsky, R., Tennenholtz, M.: Approximately optimal mechanism design via differential privacy. In: *Innovations in Theoretical Computer Science 2012*, Cambridge, MA, USA, January 8–10, 2012, pp. 203–213 (2012)
28. Penna, P., Ventre, C.: Collusion-resistant mechanisms with verification yielding optimal solutions. *ACM Trans. Comput. Theory* **4**(2), 6 (2012)
29. Penna, P., Ventre, C.: Optimal collusion-resistant mechanisms with verification. *Games Econ. Behav.* **86**, 491–509 (2014)
30. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. *ACM Trans. Econ. Comput.* **1**(4), 18 (2013)
31. Sher, I., Vohra, R.: Price discrimination through communication. *Theor. Econ.* **10**, 597–648 (2015)
32. Singh, N., Wittman, D.: Implementation with partial verification. *Rev. Econ. Des.* **6**(1), 63–84 (2001)
33. Vohra, R.V.: *Mechanism Design: A Linear Programming Approach*. Cambridge University Press, Cambridge (2011)