On the Limitations of Greedy Mechanism Design for Truthful Combinatorial Auctions

Allan Borodin[1] and Brendan Lucier[1]

[1] Dept of Computer Science,University of Toronto, bor@cs.toronto.edu.
[2] Dept of Computer Science,University of Toronto, blucier@cs.toronto.edu.

Abstract. We study mechanisms for the combinatorial auction (CA) problem, in which m objects are sold to rational agents and the goal is to maximize social welfare. Of particular interest is the special case of s-CAs, where agents are interested in sets of size at most s, for which a simple greedy algorithm obtains an $s+1$ approximation but no deterministic truthful mechanism is known to attain an approximation ratio better than $O(m/\sqrt{\log m})$. We view this as an extreme gap not only between the power of greedy auctions and truthful greedy auctions, but also as an apparent gap between the known power of truthful and non-truthful deterministic algorithms. We associate the notion of greediness with a broad class of algorithms, known as priority algorithms, which encapsulates many natural auction methods. This motivates us to ask: how well can a truthful greedy algorithm approximate the optimal social welfare for CA problems? We show that no truthful greedy priority algorithm can obtain an approximation to the CA problem that is sublinear in m, even for s-CAs with $s \geq 2$. We conclude that any truthful combinatorial auction mechanism with non-trivial approximation factor must fall outside the scope of many natural auction methods.

1 Introduction

The field of algorithmic mechanism design attempts to bridge the competing demands of agent selfishness and computational constraints. The difficulty in such a setting is that agents may lie about their inputs in order to obtain a more desirable outcome. It is often possible to circumvent this obstacle by using payments to elicit truthful responses. Indeed, if the goal of the algorithm is to maximize the total welfare of all agents, the well-known VCG mechanism does precisely that: each agent maximizes his utility by reporting truthfully. However, the VCG mechanism requires that the underlying optimization problem be solved exactly, and is therefore ill-suited for computationally intractable problems. Indeed, approximation algorithms do not, in general, result in truthful mechanisms when coupled with the VCG construction. Determining the power of truthful *approximation* mechanisms to maximize social welfare is a fundamental problem in algorithmic mechanism design.

The combinatorial auction (CA) problem holds a position at the center of this conflict between truthfulness and approximability. In this problem, m objects are to be distributed among n bidders. Each bidder holds a private value for each possible *subset* of the objects. The generality of this problem models situations in which the objects for sale can exhibit complementarities; for example, an agent's value for a pair of shoes can be much greater than twice his value for only a left shoe or a right shoe in isolation. Combinatorial auctions have arisen in pracitice for the sale of airport landing schedules [48], FCC spectrum auctions [16], and others; see [17] for an overview. The CA problem, also known as the winner determination problem for CAs, is to determine, given the agents' valuation functions (either directly or via oracle access), the allocation of objects that maximizes the overall social welfare.

Without strategic considerations, one can obtain an $O(\min\{n, \sqrt{m}\})$ approximation for CAs with n bidders and m objects with a conceptually simple (albeit not obvious) greedy algorithm [40], and this is the best possible under standard complexity assumptions [31, 49]. However, no deterministic truthful mechanism is known to obtain an approximation ratio better than $O(\frac{m}{\sqrt{\log m}})$ for the general problem [32]. This is true even for the special case where each bidder is interested only in sets of size at most some constant $s \geq 2$ (the s-CA problem), where the standard greedy algorithm obtains an $s+1$ approximation. Whether these gaps are essential for the CA problem, or whether there is some universal reduction by which approximation algorithms for the CA problem can be made truthful without heavy loss in performance, is a central open question that has received significant attention over the past decade [26, 36, 40, 45, 47].

It is known that there do exist problems for which deterministic truthful mechanisms must achieve significantly worse approximations to the social welfare than their non-truthful counterparts. Indeed, a lower bound for the related combinatorial public project problem [47] shows that there is a large asymptotic gap separating approximation by deterministic algorithms and by deterministic truthful mechanisms in general allocation problems. However, for the general CA problem, the question of whether such an essential gap exists remains open. Currently, the only lower bounds known for the general CA problem are limited to max-in-range (MIR) algorithms [22, 13]. While many known truthful CA algorithms are MIR, the possibility yet remains that non-MIR algorithms could be used to bridge the gap between truthful and non-truthful CA design.

Significant work has focused on particular restricted cases of the combinatorial auction problem, such as submodular auctions [25, 24, 34, 39, 20], and on alternative solution concepts such as randomized notions of truthfulness [18, 38, 21, 2, 23, 19], truthfulness in Bayesian settings [30, 29, 8], and performance at (non-truthful) equilibrium [43, 42, 41]. Nevertheless, the original problem yet stands as a core demonstration of the limits of our understanding of truthful approximation algorithms. A resolution would also be of practical interest, as any new insights would likely contribute, even if only indirectly, to the growing interest in robust combinatorial auction mechanisms with desirable game-theoretic properties. Indeed, many mechanisms used in practice today are based upon iterative price-determination methods which appear to work well empirically, but do not have the theoretical soundness of single-item auction methods. This situation is due at least in part to the difficulty in resolving the computational and game-theoretic issues in the CA problem from a purely theoretical perspective.

The hope, when studying the combinatorial auction problem, would be to find a natural and truthful approximation mechanism, of the flavour of the well-known Vickrey auction for a single object. This qualifier "natural" is highly subjective, but nevertheless important; it is crucial that agents understand any auction that they are participating in, even if it is truthful. Indeed, the inscrutable nature of the VCG auction has been cited as one reason why it is rarely used in practice, even in settings where optimal outcomes can be computed efficiently; it is important that agents be able to quickly determine which bids would "win" in a given auction instance [3]. We are therefore motivated to ask the following loosely-defined question: can any "natural" auction that proceeds by ranking bids in some manner, and allocating to the agents with the "best" bids, be simultaneously truthful and achieve a good approximation to the social welfare?

One may be tempted to respond to this question negatively. However, we would argue that this is not immediately clear. Indeed, many different auction methods may fit the above description. For instance, a truthful auction due to Bartal et al. [7] for the multi-unit combinatorial auction problem is a primal-dual algorithm that proceeds by iteratively constructing a price vector, and can be viewed as resolving bids in a (specially-tailored and adaptive) greedy manner. This approach is particularly appealing from a practical standpoint, as it mirrors ascending price vector methods currently in use. Is it possible that such methodologies, extended further, might lead to a similar greedy-like truthful algorithm for the more general combinatorial auction problem?

Our goal in this work will be to develop lower bounds for truthful CA mechanisms that satisfy our notion of a "natural" auction alluded to above. We ask: can any truthful *greedy* algorithm obtain an approximation ratio better than $O(\frac{m}{\sqrt{\log(m)}})$? Our specific interest in greedy algorithms is motivated threefold. First, most known examples of truthful, non-MIR algorithms for combinatorial auction problems apply greedy methods [4, 7, 12, 14, 35, 39, 40, 45]; indeed, greedy algorithms embody the conceptual monotonicity properties generally associated with truthfulness, and are thus natural candidates for truthful mechanism construction. Second, greedy algorithms are known to obtain asymptotically tight approximation bounds for many CA problems despite their simplicity. Finally, and perhaps most importantly, many auctions used in practice apply greedy methods, despite the fact that they may not be incentive compatible (e.g. the generalized second price auction for adwords [27]). That is, simple mechanisms (and in particular greedy mechanisms) seem to be good candidates for auctions due to other considerations beyond truthfulness, such as ease of public understanding and perceived fairness.

We use the term "greedy algorithm" to refer to any of a large class of algorithms known as *priority algorithms* [11]. The class of priority algorithms captures a general notion of greedy algorithm behaviour. Informally speaking, a priority algorithm proceeds by ranking bids according to some (possibly adaptive) quality score; winning a certain bundle then requires that one's bid be superior (in terms of the ranking) to the conflicting bids with which it competes. Priority algorithms include, for example, many well-known primal-dual algorithms, as well as other greedy algorithms with adaptive and non-trivial selection rules. Moreover, this class is *independent of computational constraints* and also independent of the manner in which valuation functions are accessed. In particular, our results apply to algorithms in the demand query model and the general query model, as well as to auctions in which bids are explicitly represented. Roughly speaking, a priority algorithm has some notion of what constitutes the "best" bid in any given auction instance; the auction finds this bid, satisfies it, then iteratively resolves the reduced auction problem with fewer objects (possibly with an adaptive notion of the "best" bid). For example, the previously mentioned truthful algorithm for multi-unit auctions due to Bartal et al. [7] that updates a price vector while iteratively satisfying agent demands falls into this framework. Our main result demonstrates that if a truthful auction for an s-CA proceeds in this way, then it cannot perform much better than the naive algorithm that allocates all objects to a single bidder. The gap described in our result is extreme: for $s = 2$, the standard (but non-truthful) greedy algorithm is a 3-approximation for the s-CA problem, but no truthful greedy algorithm can obtain a sublinear approximation bound.

We also consider the combinatorial auction problem for submodular bidders (SMCA), a very well-studied special case of the general CA problem. We study a class of greedy algorithms that is especially well-suited to the SMCA problem. Such algorithms consider the objects of the auction one at a time and greedily assign them to bidders to maximize marginal utilities. It was shown in [39] that any such algorithm[3] attains a 2-approximation to the SMCA problem, but that not all are incentive compatible. We show that, in fact, no such algorithm can be incentive compatible.

1.1 Related Work

There have been many developments in the restricted case of CAs with single-minded bidders. Following the Lehmann et al. [40] truthful greedy mechanism for single-minded CAs, Mu'alem and Nisan [45] showed that any *monotone* greedy algorithm for single-minded bidders is truthful, and outlined various techniques for combining approximation algorithms while retaining truthfulness. This led to the development of many other truthful algorithms in single-minded settings [5, 12] and additional construction techniques, such as the iterative greedy packing of loser-independent algorithms due to Chekuri and Gamzu [14].

Less is known in the setting of general bidder valuations. The best-known truthful deterministic mechanism for the general CA problem proceeds by dividing the objects arbitrarily into $O(\log m)$ equal-sized indivisible bundles, then allocating those bundles optimally; this achieves an approximation ratio of $O(m/\log m)$

[3] The degree of freedom in this class of algorithms is the order in which the objects are considered.

[32]. For the special case of multi-unit CAs, when there are $B \geq 3$ copies of each object, Bartal et al. [7] give a greedy algorithm that obtains an $O(Bm^{\frac{1}{B-2}})$ approximation. Lavi and Swamy [38] give a general method for constructing randomized mechanisms that are truthful in expectation, meaning that agents maximize their expected utility by declaring truthfully. Their construction generates a k-approximate mechanism from an LP for which there is an algorithm that verifies a k-integrality gap, and in particular they obtain an $O(\sqrt{m})$ approximation for the general CA problem. In the applications they discuss, these verifiers take the form of greedy algorithms, which play a prominant role in the final mechanisms. Dobzinski, Nisan, and Schapira [24] construct a universally truthful randomized $O(\sqrt{m})$-approximate mechanism for the CA problem via sampling.

A significant line of research aims to give lower bounds on the approximating power of deterministic truthful algorithms for CAs. Lehmann, Mu'alem, and Nisan [36] show that any truthful CA mechanism that uses a suitable bidding language, is unanimity-respecting, and satisfies an independence of irrelevant alternatives property (IIA) cannot attain a polynomial approximation ratio. It has also been shown that, roughly speaking, any truthful polytime subadditive combinatorial auction mechanism with an approximation factor better than 2 cannot satisfy the natural property of being *stable*[4] [26]. Dobzinski and Nisan showed that no max-in-range algorithm can obtain an approximation ratio better than $\Omega(\sqrt{m})$ with polynomial communication between agents and the mechanism [22]. This was later extended to show that no max-in-range algorithm can obtain an approximation ratio better than $\Omega(\sqrt{m})$ even when agents have succinctly-representable valuations (i.e., budget-constrained additive valuations) [13]. These lower bounds are incomparable to our own, as priority algorithms need not be MIR, stable, unanimity-respecting, or satisfy IIA[5].

There has been extensive work studying the power of truthful mechanisms for restricted forms of combinatorial auctions, such as submodular auctions [25, 24, 34, 39, 20, 23]. Of particular relevance to our work is the recent work of Dobzinski [20] that establishes a large gap between the power of randomized algorithms and universally truthful randomized mechanisms for the submodular CA problem, in the value oracle query model. Specifically, a universally truthful mechanism requires exponentially many queries to obtain approximation ratio $O(m^{\frac{1}{2}-\epsilon})$; this bound closely matches the $O(m^{\frac{1}{2}})$ approximation attainable by a deterministic truthful mechanism [39].

Another line of work gives lower bounds for greedy algorithms without truthfulness restrictions. Gonen and Lehmann [28] showed that no algorithm that greedily accepts bids for sets can guarantee an approximation better than $\sqrt{m}$

[4] In a stable mechanism, no player can alter the outcome (i.e. by changing his declaration) without causing his own allocated set to change.

[5] The notion of IIA has been associated with priority algorithms, but in a different context than in [36]. In mechanism design, IIA is a property of the mapping between input valuations and output allocations, whereas for priority algorithms the term IIA describes restrictions on the order in which input items can be considered.

for the general CA problem. More generally, Krysta [35] showed that no oblivious greedy algorithm (in our terminology: fixed order greedy priority algorithm) obtains approximation ratio better than $\sqrt{m}$. In contrast, we consider the even more general class of all priority algorithms but restrict them to be incentive compatible.

The class of priority algorithms is loosely related to the notion of online algorithms. Mechanism design has been studied in a number of online settings, and lower bounds are known for the performance of truthful algorithms in these settings [37, 44]. The critical difference between these results and our lower bounds is that a priority algorithm has control over the order in which input items are considered, whereas in an online setting this order is chosen adversarily.

In contrast to the negative results of this paper, (non-truthful) greedy algorithms can provide good approximations when rational agents are assumed to bid at Nash equilibrium. In particular, there is a greedy combinatorial auction for submodular agents that obtains a 2-approximation at any Bayes-Nash equilibrium [15], and a similar auction method obtains a 2-approximation at equilibrium and a $2 \log m$-approximation at Bayes-Nash equilibrium for subadditive bidders [9]. The greedy GSP auction for internet advertising has been shown to obtain a 1.6-approximation at pure Nash equilibrium [41] and a 3.1-approximation at Bayes-Nash equilibrium. It is also known that, in a wide variety of contexts, c-approximate greedy algorithms for combinatorial allocation problems can be converted into mechanisms whose Bayes-Nash equilibria yield $c(1 + o(1))$ approximations [43, 42].

2 Defintions and Preliminary Results

2.1 Combinatorial Auctions

A *combinatorial auction* consists of n bidders and a set M of m objects. Each bidder i has a value for each subset of objects $S \subseteq M$, described by a valuation function $v_i : 2^M \to \mathbb{R}$ which we call the *type* of agent i. We assume each v_i is monotone and normalized so that $v_i(\emptyset) = 0$. We denote by V_i the space of all possible valuation functions for agent i, and $V = V_1 \times V_2 \times \cdots \times V_n$. We write $\mathbf{v}$ for a profile of n valuation functions, one per agent, and $\mathbf{v}_{-i} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$, so that $\mathbf{v} = (v_i, v_{-i})$.

A valuation function v is *single-minded* if there exists a set $S \subseteq M$ and a value $x \geq 0$ such that, for all $T \subseteq M$, $v(T) = x$ if $S \subseteq T$ and 0 otherwise. A valuation function v is *k-minded* if it is the maximum of k single-minded functions. That is, there exist k sets $S_1, \ldots, S_k$ such that for all subsets $T \subseteq M$ we have $v(T) = \max\{v(S_i) | S_i \subseteq T\}$. An *additive* valuation function v is specified by m values $x_1, \ldots, x_m \in \mathbb{R}_{\geq 0}$ so that $v(T) = \sum_{a_i \in T} x_i$. A valuation function v is *submodular* if it satisfies $v(T) + v(S) \geq v(S \cup T) + v(S \cap T)$ for all $S, T \subseteq M$.

A *direct revelation mechanism* (or just *mechanism*) $\mathcal{M} = (G, P)$ consists of an *allocation algorithm* G and a *payment algorithm* P. We think of $\mathcal{M}$ as eliciting a valuation profile from the agents, and then determining an assignment of objects and payments to be made. Crucially, the agents are assumed to be rational and may misrepresent their valuations to the mechanism but only in order to maximize their utilities. We therefore distinguish between the values reported to the mechanism from the agents' true values: we shall let $\mathbf{d}$ denote a profile of declared valuations and let $\mathbf{t}$ denote a truthful valuation profile. Given declared valuation profile $\mathbf{d}$, $G(\mathbf{d})$ returns an allocation of objects to bidders, and $P(\mathbf{d})$ returns the payment extracted from each agent. For each agent i we write $G_i(\mathbf{d})$ and $P_i(\mathbf{d})$ for the set given to and payment extracted from i.

The *social welfare* obtained by G on declaration $\mathbf{d}$ given truthful valuation $\mathbf{t}$ is $SW_G(\mathbf{d}, \mathbf{t}) = \sum_{i \in N} t_i(G_i(\mathbf{d}))$. The *optimal social welfare*, SW_{opt}, is the maximum of $\sum_{i \in N} t_i(S_i)$ over all valid allocations $(S_1, \ldots, S_n)$. Algorithm G is a c-approximation if $SW_G(\mathbf{t}, \mathbf{t}) \geq \frac{1}{c} SW_{opt}$ for all type profiles $\mathbf{t}$.

Fixing mechanism $\mathcal{M}$ and type profile $\mathbf{t}$, the *utility* of bidder i given declaration $\mathbf{d}$ is $u_i(\mathbf{d}) = t_i(G_i(\mathbf{d})) - P_i(\mathbf{d})$. Mechanism $\mathcal{M}$ is said to be *truthful in dominant strategies* (or just *truthful*) if, for every type profile $\mathbf{t}$, agent i, and declaration profile $\mathbf{d}$, $u_i(t_i, \mathbf{d}_{-i}) \geq u_i(\mathbf{d})$. That is, agent i maximizes his utility by declaring his type, regardless of the declarations of the other agents. We say that G is truthful (or *incentive compatible*) if there exists a payment function P such that the mechanism (G, P) is truthful.

2.2 Critical Prices

An allocation algorithm G defines *critical prices*, $p_i(S, \mathbf{v}_{-i})$, for any agent i and set S. The price $p_i(S, \mathbf{v}_{-i})$ is the minimum amount that agent i could bid on set S and still win S assuminng the other agents bid according to $\mathbf{v}_{-i}$. That is,

$$p_i(S, \mathbf{v}_{-i}) = \inf\{v : \exists \mathbf{d} \text{ such that } d_i(S) = v \text{ and } G_i(\mathbf{d}) = S\}.$$

From Bartal, Gonen and Nisan [7], we have the following characterization of truthful mechanisms for combinatorial auctions, which we will use throughout.

Theorem 1. *A mechanism $\mathcal{M} = (G, P)$ is truthful if and only if, for every bidder i and every vector of bids of the other bidders $\mathbf{d}_{-i}$,*

1. $P_i(\mathbf{d}) = p_i(G_i(\mathbf{d}), \mathbf{d}_{-i})$ *and*
2. $G_i(\mathbf{d}) \in \arg\max_S\{d_i(S) - p_i(S, \mathbf{d}_{-i})\}$.

Theorem 1 states that a truthful mechanism must charge critical prices as its payments, and must always allocate to each agent a set that maximizes his utility subject to these prices.

Note that $p_i(S, \mathbf{d}_{-i})$ need not be finite. If $p_i(S, \mathbf{d}_{-i}) = \infty$, then a mechanism will simply not allocate S to bidder i for any reported valuation d_i. In addition, as we now show, one can assume without loss of generality that critical prices are monotone in the allocated set S.

Claim. Suppose that a mechanism satisfies the conditions of Theorem 1 (i.e. the mechanism is truthful). Then one can assume without loss of generality that for all $i \in N$, all $\mathbf{d}_{-j} \in V_{-i}$, and all $S \subseteq T \subseteq M$, $p_i(S, \mathbf{d}_{-i}) \leq p_i(T, \mathbf{d}_{-i})$.

Proof. Suppose $\mathcal{M} = (G, P)$ is an incentive compatible mechanism; then $\mathcal{M}$ satisfies the critical pricing property, say with prices p_i. We will construct a new set of critical prices $p_i{}'$ that satisfies the conditions of Theorem 1 while also satisfying the monotonicity conditions of our claim.

For all $i \in N$, $\mathbf{d}_{-i} \in V_{-i}$, and $S \subseteq M$, define $p_i{}'(S, \mathbf{d}_{-i})$ by

$$p_i{}'(S, \mathbf{d}_{-i}) = \min\{p_i(T, \mathbf{d}_{-i}) | T \supseteq S\}.$$

Notice that $p_i{}'(S, \mathbf{d}_{-i}) \leq p_i{}'(T, \mathbf{d}_{-i})$ whenever $S \subseteq T$. Furthermore, $p_i{}'(S, \mathbf{d}_{-i}) \leq p_i(S, \mathbf{d}_{-i})$ for all $S \subseteq M$.

We claim that $p_i{}'$ satisfies the conditions of Theorem 1 for the mechanism $\mathcal{M}$. Choose any $i \in N$ and $\mathbf{d} \in V$ and suppose that $G_i(\mathbf{d}) = \tilde{S}$. Then, by the critical pricing property (for prices p_i), $d_i(\tilde{S}, \mathbf{d}_{-i}) - p_i(\tilde{S}, \mathbf{d}_{-i}) \geq d_i(T, \mathbf{d}_{-i}) - p_i(T, \mathbf{d}_{-i})$ for all $T \subseteq M$, and furthermore $P_i(\mathbf{d}) = p_i(\tilde{S}, \mathbf{d})$. We need to show that if the mechanism sets $P_i(\mathbf{d}) = p_i{}'(\tilde{S}, \mathbf{d})$ then $d_i(\tilde{S}, \mathbf{d}_{-i}) - p_i{}'(\tilde{S}, \mathbf{d}_{-i}) \geq d_i(T, \mathbf{d}_{-i}) - p_i{}'(T, \mathbf{d}_{-i})$ for all $T \subseteq M$. We derive the following for all $T \subseteq M$:

$$\begin{aligned}
d_i(T, \mathbf{d}_{-i}) - p_i{}'(T, \mathbf{d}_{-i}) &= d_i(T, \mathbf{d}_{-i}) - p_i(T', \mathbf{d}_{-i}) \text{ for some } T' \supseteq T \\
&\leq d_i(T', \mathbf{d}_{-i}) - p_i(T', \mathbf{d}_{-i}) \text{ by monotonicty of declarations} \\
&\leq d_i(\tilde{S}, \mathbf{d}_{-i}) - p_i(\tilde{S}, \mathbf{d}_{-i}) \text{ since the } p_i \text{ are critical prices} \\
&\leq d_i(\tilde{S}, \mathbf{d}_{-i}) - p_i{}'(\tilde{S}, \mathbf{d}_{-i}) \text{ by the definition of } p_i{}'
\end{aligned}$$

We therefore have that the $p_i{}'$ are critical prices for the mechanism, as required. □

2.3 Priority Algorithms

In this section we review the priority algorithm framework [11] and discuss how it can be applied to the CA problem. We view an *input instance* to an algorithm as a subset of *input items* from a known input space $\mathcal{I}$. Note that $\mathcal{I}$ depends on the problem being considered, and is the set of *all possible* input items: an input instance is a finite subset I of $\mathcal{I}$. The problem definition may place restrictions on the input: an input instance $I \subseteq \mathcal{I}$ is *valid* if it satisfies all such restrictions. For example, in the CA problem, we would not allow having an agent who values a subset $S' \subset S$ more than the set S. The output of the algorithm is a decision made for each input item in the input instance. For example, these decisions may be of the form "accept/reject", allocate set S to agent i, etc. The problem may place restrictions on the nature of the decisions made by the algorithm; we say that the output of the algorithm is *valid* if it satisfies all such restrictions. A *priority algorithm* is then any algorithm of the following form:

ADAPTIVE PRIORITY
Input: A set I of items, $I \subseteq \mathcal{I}$
while not empty(I)
  Ordering: Choose, without looking at I, a total ordering $\mathcal{T}$ over $\mathcal{I}$

next ← first item in I according to ordering $\mathcal{T}$
Decision: make an irrevocable decision for item *next*
remove *next* from I; remove from $\mathcal{I}$ any items preceding *next* in $\mathcal{T}$
end while

We emphasize the importance of the ordering step in this framework: an adaptive priority algorithm is free to choose *any* ordering over the space of possible input items, and can change this ordering adaptively after each input item is considered. Once an item is processed, the algorithm is not permitted to modify its decision. On each iteration a priority algorithm learns what (higher-priority) items are *not* in the input. A special case of (adaptive) priority algorithms are *fixed order* priority algorithms in which one fixed ordering is chosen before the while loop (i.e. the "ordering" and "while" statements are interchanged). Our inapproximation results for truthful CAs will hold for the more general class of adaptive priority algorithms although many greedy CA algorithms are fixed order.

Admittedly, the term "greedy" implies a more opportunistic aspect than is apparent in the definition of priority algorithms. Indeed, we view priority algorithms more generally as "greedy-like" or "myopic". A *greedy* priority algorithm satisfies an additional property: the choice made for each input item must optimize the objective of the algorithm as though that item were the last item in the input. We note that many greedy CA algorithms are fixed order greedy priority algorithms.

3 Truthful Priority Algorithms

As noted above, Lehmann, O'Callahan and Shoham [40] show that a greedy $O(\sqrt{m})$-approximation algorithm [6] for combinatorial auctions can be made truthful (using critical pricing) for single-minded bidders, but is not incentive compatible for the more general CA problem. Our high-level goal is to prove that this is a general phenomenon common to all priority algorithms. In order to apply the concept of priority algorithms we must define the set $\mathcal{I}$ of possible input items and the nature of decisions to be made. We consider two natural input formulations: sets as items, and bidders as items. We assume that n, the number of bidders, and m, the number of objects, are known to the mechanism and let $k = \min\{m, n\}$.

3.1 Sets as Items

In our primary model, we view an input instance to the combinatorial auction problem as a list of set-value pairs for each bidder. An item is a tuple (i, S, t), $i \in N$, $S \subseteq M$, and $t \in \mathbb{R}_{\geq 0}$. A valid input instance $I \subset \mathcal{I}$ contains at most one tuple $(i, S, v_i(S))$ for each $i \in N$ and $S \subseteq M$ and for every pair of tuples (i, S, v) and (i', S', v') in I such that $i = i'$ and $S \subseteq S'$, it must be that $v \leq v'$.

[6] The Lehmann et al algorithm will satisfy all models discussed in section 3.

We note that since a valid input instance may contain an exponential number of items, this model applies most directly to algorithms that use oracles to query input valuations, such as demand oracles[7], but it can also apply to succinctly represented valuation functions.[8]

The decision to be made for item (i, S, t) is whether or not the objects in S should be added to any objects already allocated to bidder i. For example, an algorithm may consider item (i, S_1, t_1) and decide to allocate S_1 to bidder i, then later consider another item (i, S_2, t_2) (where S_2 and S_1 are not necessarily disjoint) and, if feasible, decide to change bidder i's allocation to $S_1 \cup S_2$.

A *greedy algorithm* in the sets as items model must accept any feasible, profitable item (i, S, t) it considers.[9] Our main result is a lower bound on the approximation ratio achievable by a truthful greedy algorithm in the sets as items model. Theorem 2 implies a severe separation between the power of greedy algorithms and the power of truthful greedy algorithms. A simple greedy algorithm obtains a 3-approximation for the 2-CA problem, yet no truthful greedy priority algorithm (indeed, any algorithm that irrevocably satisfies bids based on a notion of priority) can obtain even a sublinear approximation.

Theorem 2. *Suppose A is an incentive compatible greedy priority algorithm that uses sets as items. Then A cannot approximate the optimal social welfare by a factor of $\frac{(1-\delta)k}{2}$ for any $\delta > 0$. This result also applies to the special case of the 2-CA problem, in which each desired set has size at most 2.*

Before beginning the proof, consider the following intuition as to why such an algorithm A cannot exist. Suppose some bidder i has a very large value for each of two singletons. Our algorithm A would surely want to allocate one of these singletons to this bidder. Since A is greedy, it must do so without first considering the (smaller) values held by other bidders for sets containing those singletons. However, if A is truthful, then by Theorem 1 it must also maximize utility for agent i. The algorithm must therefore allocate the singleton which has the smaller critical price. This implies that the relationship between the prices for these singletons must be independent of their value to other bidders! This allows us to show that algorithm A must have poor performance, since a singleton desired at a high value by many players must have a higher price

[7] It is tempting to assume that this model is equivalent to a value query model, where the mechanism queries bidders for their values for given sets. The priority algorithm model is actually more general, as the mechanism is free to choose an arbitrary ordering over the space of possible set/value combinations. In particular, the mechanism could order the set/value pairs by the utility they would generate under a given set of additive prices, simulating a demand query oracle.

[8] That is, by assigning priority only to those tuples appearing in a given representation.

[9] That is, when considering a bid (i, S, t), a greedy algorithm must allocate S to agent i if no objects in S have already been allocated to another bidder, and $d_i(S_1 \cup S) > d_i(S_1)$. In our proof of Theorem 2, it will always be the case that $S_1 = \emptyset$ (i.e. no items have already been allocated to agent i), so that the greedy assumption is simplified as follows: when considering a bid (i, S, t), a greedy algorithm must allocate S to agent i if $t > 0$ and no objects in S have already been allocated to another bidder.

than a singleton not desired by any other players, in order to guarantee a good approximation ratio.

Proof. Choose $\delta > 0$ and suppose A obtains a bounded approximation ratio. For each $i \in N$, let V_{-i}^{+} be the set of valuations with the property that $v_\ell(S) > 0$ for all $\ell \neq i$ and all non-empty $S \subseteq M$. The heart of our proof is the following claim, which shows that the relationship between critical prices for singletons for one bidder is independent of the valuations of other bidders. Recall that $p_i(S, \mathbf{d}_{-i})$ is the critical price for set S for bidder i, given $\mathbf{d}_{-i}$.

Lemma 1. *For all $i \in N$, and for all $a, b \in M$, either $p_i(\{a\}, \mathbf{d}_{-i}) \geq p_i(\{b\}, \mathbf{d}_{-i})$ for all $\mathbf{d}_{-i} \in V_{-i}^{+}$, or $p_i(\{a\}, \mathbf{d}_{-i}) \leq p_i(\{b\}, \mathbf{d}_{-i})$ for all $\mathbf{d}_{-i} \in V_{-i}^{+}$. This is true even when agents desire sets of size at most 2.*

Proof. Choose $i \in N$, $a, b \in M$, and $\mathbf{d}_{-i}, \mathbf{d}_{-i}' \in V_{-i}^{+}$. Suppose for contradiction that $p_i(\{a\}, \mathbf{d}_{-i}) > p_i(\{b\}, \mathbf{d}_{-i})$ but $p_i(\{b\}, \mathbf{d}_{-i}') > p_i(\{a\}, \mathbf{d}_{-i}')$. We will consider a number of possible valuations to be declared by our bidders.

Let v^* be the maximum value assigned to any set by any player in $\mathbf{d}_{-i}$ or $\mathbf{d}_{-i}'$. Then note that the maximum social welfare that can be obtained is $(k-1)v^*$ if bidder i does not participate and other bidders declare values $\mathbf{d}_{-i}$ or $\mathbf{d}_{-i}'$. Let $x = k^2 v^*$. We will define various different possible valuation functions for bidder i: f, h, and g_c for all $c \in M$.

$$f(S) = \begin{cases} x & \text{if } a \in S \\ x & \text{if } b \in S \\ 0 & \text{otherwise.} \end{cases} \qquad g_c(S) = \begin{cases} \epsilon & \text{if } a \in S,\ c \notin S \\ \epsilon & \text{if } b \in S \\ x & \text{if } \{a, c\} \subseteq S \\ 0 & \text{otherwise.} \end{cases}$$

$$h(S) = \begin{cases} \epsilon & \text{if } a \in S \\ \epsilon & \text{if } b \in S \\ 0 & \text{otherwise.} \end{cases}$$

Note that each of these valuation profiles can be interpreted as a profile in which the agent desires sets of size at most 2. Note also that g_a and g_b are well-defined: the former assigns value x to any set containing a, and the latter assigns value x to any set containing both a and b.

We are now ready to discuss the behaviour of algorithm A. Consider the subset $\mathcal{I}_1 \subset \mathcal{I}$ that contains the following input items: $(i, S, f(S))$ and $(i, S, h(S))$ for every $S \subseteq M$; $(i, S, g_c(S))$ for all $c \in M$ and $S \subseteq M$; and $(j, S, d_j(S))$, $(j, S, d_j'(S))$, (j, S, ϵ), and (j, S, v^*) for all $j \neq i$ and $S \subseteq M$. In other words, $\mathcal{I}_1$ contains all of the input items consistent with the valuation functions we defined above, plus input items (j, S, ϵ) and (j, S, v^*) for each set S and each bidder $j \neq i$.

We know that if A is a priority algorithm, then it must have some initial ordering over $\mathcal{I}$, and hence over $\mathcal{I}_1$. Consider the first item in $\mathcal{I}_1$ under this ordering. We consider different cases for the nature of this item.

Case 1: (j, S, t)**,** $j \neq i$. Then $t \in \{d_j(S), d_j{}'(S), \epsilon, v^*\}$ and hence $t > 0$. Choose any $c \in S$. Let I_1 be a valid input instance consisting of items from $\mathcal{I}_1$, such that $(j, S, t) \in I_1$ and I_1 is consistent with agent i having valuation g_c. Note that such an I_1 always exists; for example, if $t = d_j(S)$ we could set I_1 to be consistent with each agent $\ell \neq i$ having valuation d_ℓ. Then $I_1 \subseteq \mathcal{I}_1$ and $(j, S, t) \in I_1$, so item (j, S, t) will be considered first by algorithm A on input I_1.

Since A is greedy, A will allocate set S to bidder j. Then it must be that, in the final allocation, bidder i is not allocated any set containing c. Thus, from the definition of g_c, bidder i obtains a value of at most ϵ. Furthermore, all other bidders can obtain a total welfare of at most $(k-1)v^*$, for a total social welfare of at most $(k-1)v^*+\epsilon$. On the other hand, a total of at least $x = k^2v^*$ is possible by allocating $\{a, c\}$ to bidder i. Then as long as $\epsilon < v^*$ the approximation ratio obtained by A is at least k, a contradiction.

The other cases for t are handled similarly.

Case 2: (i, S, x)**,** $a \in S$ **or** $b \in S$. By symmetry we can assume $a \in S$. Consider the input instance I_2 in which bidder i declares valuation f, and every other bidder $j \neq i$ declares valuation d_j. Then $f(S) = x$, so $(i, S, x) \in I_2 \subseteq \mathcal{I}_1$, and therefore A will consider item (i, S, x) first on input I_2. Since $x > 0$ and A is greedy, the algorithm will assign set S to bidder i.

Suppose that in the final allocation, bidder i is allocated some set $T \supseteq S$. Then since $a \in T$, we know that $p_i(T, \mathbf{d}_{-i}) \geq p_i(\{a\}, \mathbf{d}_{-i}) > p_i(\{b\}, \mathbf{d}_{-i})$. But note $f(T) = f(\{a\}) = x$, so that $f(T) - p_i(T, \mathbf{d}_{-i}) < f(\{b\}) - p_i(\{b\}, \mathbf{d}_{-i})$. In other words, A does not maximize the utility of player i. By Theorem 1, A is not incentive compatible, a contradiction.

Case 3: (i, S, ϵ)**,** $a \in S$ **or** $b \in S$. By symmetry we can assume $a \in S$. Consider the input instance I_3 in which bidder i declares valuation h, and every other bidder $j \neq i$ declares valuation d_j. Then $(i, S, \epsilon) \in I_3 \subseteq \mathcal{I}_1$, so A will consider item (i, S, ϵ) first on input I_3. From this point, we obtain a contradiction in precisely the same way as in Case 2.

Case 4: (i, S, t)**,** $a \notin S$ **and** $b \notin S$. Then from the definitions of f, g_c, and h, we must have $t = 0$. Thus when processing this item, A is free to allocate S to bidder i or not. If A does not allocate S to i, then we will consider the *next* item considered by the algorithm A, and repeat our case analysis. The case analysis proceeds in the same way, since no objects would have been allocated. This process must terminate, as algorithm A must eventually consider some set S for agent i that contains either a or b, or (reasoning as above) some set for agent $j \neq i$.

Suppose, on the other hand, that A does allocate S to i. Then consider the input instance I_4 in which bidder i declares valuation h and all other bidders declare the following valuation f_S:

$$f_S(T) = \begin{cases} v^* & \text{if } S \subseteq T \\ \epsilon & \text{otherwise.} \end{cases}$$

We note that valuation f_S defines the value of any set to be either ϵ or v^*, so in particular $I_3 \subseteq \mathcal{I}_1$. Since $(i, S, 0) \in I_3$, this item will be considered first by A on input I_3, and S will be allocated to player i. But then in the final allocation each other bidder can obtain a welfare of at most ϵ, for a total welfare of at most $k\epsilon$. On the other hand, a welfare of v^* was possible by allocating S to any bidder other than bidder i. Thus, if we choose $\epsilon < v^*/k^2$ we conclude that A has an approximation ratio of at least k, a contradiction.

We have shown that every case leads to a contradiction, completing the proof of Lemma 1. $\square$

We can think of Lemma 1 as defining, for each $i \in N$, an ordering over the elements of M. For each $i \in N$ and $a, b \in M$, write $a \preceq_i b$ to mean $p_i(a, \mathbf{d}_{-i}) \leq p_i(b, \mathbf{d}_{-i})$ for all $\mathbf{d}_{-i} \in V_{-1}^{+}$. For all $i \in N$ and $a \in M$, define $T_i(a) = \{a_j : a \preceq_i a_j\}$. That is, $T_i(a)$ is the set of objects that have higher price than a for agent i. Our next claim shows a strong relationship between whether a is allocated to bidder i and whether any object in $T_i(a)$ is allocated to bidder i.

Lemma 2. *Choose $a \in M$, $i \in N$, and $S \subseteq M$, and suppose $S \cap T_i(a) \neq \emptyset$. Choose some $d_i \in V_i$ and suppose that $d_i(\{a\}) > d_i(S)$. Then if $\mathbf{d}_{-i} \in V_{-i}^{+}$, bidder i cannot be allocated set S by algorithm A given input $\mathbf{d}$.*

Proof. We know that $p_i(S, \mathbf{d}_{-i}) \geq p_i(\{a_j\}, \mathbf{d}_{-i})$ for any $a_j \in S$. Thus, regardless of the choice of $\mathbf{d}_{-i}$,

$$p_i(S, \mathbf{d}_{-i}) \geq \max_{a_j \in S \cap T_i(a)} (p_i(\{a_j\}, \mathbf{d}_{-i})) \geq p_i(\{a\}, \mathbf{d}_{-i})$$

from the definition of $T_i(a)$. Since $d_i(a) > d_i(S)$, this implies that $d_i(a) - p_i(\{a\}, \mathbf{d}_{-i}) > d_i(S) - p_i(S, \mathbf{d}_{-i})$, so by Theorem 1 bidder i cannot be allocated set S, as required. $\square$

Lemma 2 is strongest when $T_i(a)$ is large; that is, when a is "small" in the ordering $\preceq_i$. We therefore wish to find an object of M that is small according to many of these orderings, simultaneously. Let $R(a) = \{i \in N : |T_i(a)| \geq k/2\}$, so $R(a)$ is the set of players for which there are at least $k/2$ objects greater than a. The next claim follows by a straightforward counting argument.

Lemma 3. *There exists $a^* \in M$ such that $|R(a^*)| \geq k/2$.*

Proof. We note that

$$\sum_{i \in N} \sum_{\substack{a \in M \\ |T_i(a)| \geq k/2}} 1 = \sum_{i \in N} (m - k/2) = n(m - k/2).$$

Rearranging order of summation, we also have

$$\sum_{i \in N} \sum_{\substack{a \in M \\ |T_i(a)| \geq k/2}} 1 = \sum_{a \in M} \sum_{\substack{i \in N \\ |T_i(a)| \geq k/2}} = \sum_{a \in M} |S(a)|.$$

We conclude that $\sum_{a \in M} |S(a)| = n(m - k/2)$, so there must exist some $a^* \in M$ such that $|S(a^*)| \geq \frac{n(m-k/2)}{m}$. We know that either $n \geq m = k$ or $m \geq n = k$; in either case we obtain $|S(a^*)| \geq \frac{n(m-k/2)}{m} \geq k/2$ as required. $\square$

We are now ready to proceed with the proof of Theorem 2. Let $a^* \in M$ be the object from Lemma 3. Let $\epsilon > 0$ be a sufficiently small value to be defined later. We now define a particular input instance to algorithm A. For each $i \in R(a^*)$, bidder i will declare the following valuation function, d_i:

$$d_i(S) = \begin{cases} 1 & \text{if } a^* \in S \\ 1 - \delta/2 & \text{if } a^* \notin S \text{ and } S \cap (T_i(a^*)) \neq \emptyset \\ \epsilon & \text{otherwise.} \end{cases}$$

Each bidder $i \notin R(a^*)$ will declare a value of ϵ for every set.

For each $i \in R(a^*)$, $d_i(a_j) \geq 1-\delta/2$ for every $a_j \in T_i(a^*)$. Since $|R(a^*)| \geq k/2$ and $|T_i(a^*)| \geq k/2$, it is possible to obtain a social welfare of at least $\frac{(1-\delta/2)k}{2}$ by allocating singletons to bidders in $R(a^*)$.

Consider the social welfare obtained by algorithm A. The algorithm can allocate object a^* to at most one bidder, say bidder i, who will obtain a social welfare of at most 1. For any bidder $\ell \in R(a^*)$, $\ell \neq i$, $d_\ell(S) = 1 - \delta/2 < 1$ for any S containing elements of $T_\ell(a^*)$ but not a^*. Thus, by Lemma 2, no bidder in $R(a^*)$ can be allocated any set S that contains an element of $T_i(a^*)$ but not a^*. Therefore every bidder other than bidder i can obtain a value of at most ϵ, for a total social welfare of at most $1 + k\epsilon$.

We conclude that algorithm A has an approximation factor no better than $\frac{k(1-\delta/2)}{2(1+k\epsilon)}$. Choosing $\epsilon < \frac{\delta}{2(1-\delta)k}$ yields an approximation ratio greater than $\frac{k(1-\delta)}{2}$, completing the proof of Theorem 2. □

We believe that the greediness assumption of Theorem 2 can be removed. As partial progress toward this goal, we show that this assumption can be removed if we restrict our attention to the following alternative input model for priority algorithms, in which an algorithm can only consider and allocate sets whose values are explicitly represented (i.e. not implied by the value of a subset).

Elementary bids as items. Consider an auction setting in which agents do not provide entire valuation functions, but rather each agent specifies a list of *desired sets* $S_1, \ldots, S_k$ and a value for each one. Moreover, each agent receives either a desired set or the empty set. This can be thought of as an auction with a succinct representation for valuation functions, in the spirit of the XOR bidding language [46]. We model such an auction as a priority algorithm by considering items to be the bids for desired sets. In such a setting, the specified set-value pairs are called *elementary bids*. We say that the priority model uses *elementary bids as items* when only elementary bids $(i, S, v(S))$ can be considered by the algorithm. For each item $(i, S, v(S))$, the decision to be made is whether or not S will be the one and only one set allocated to agent i; that is, whether or not the elementary bid for S will be "satisfied." In particular, unlike in the sets as items model, we do not permit the algorithm to build up an allocation incrementally by accepting many elementary bids from a single agent.

We now show that the greediness assumption from Theorem 2 can be removed when we consider priority algorithms in the elementary bids as items model.

Theorem 3. *Suppose A is an incentive compatible priority algorithm for the CA problem that uses elementary bids as items. Then A cannot approximate the optimal social welfare by a factor of $(1-\delta)k$ for any $\delta > 0$.*

Proof. Suppose A is a truthful adaptive priority algorithm, where the items to be considered are associated with sets. That is, an item is a tuple (i, S, t) where $d_i(S) = t$. On processing each item, the algorithm must decide whether S will be the set allocated to bidder i. Suppose for contradiction that A obtains an approximation ratio of $(1-\delta)k$ for some $\delta > 0$.

We first note that if only bidder i places bids, then $p_i(M, \mathbf{d}_{-i}) = 0$.

Let I_1 be an input instance containing items $(i, M, 1+\delta)$ and $(i, S, 1)$ for all $S \neq M$, for each $1 \leq i \leq N$. That is, each bidder has a value of 1 for each singleton and $1+\delta$ for the set of all objects. Then A must consider some input item first given input I_1; suppose the first item has corresponding bidder j. Now consider cases based on the nature of the first item.

Case 1: $(j, M, 1+\delta)$. Consider the decision made by A for this item. If A allocates M to j, then for input instance I_1 A obtains a social welfare of $1+\delta$, whereas the optimal welfare is k. Thus A has an approximation ratio no better than $(1+\delta)^{-1}k > (1-\delta)k$, a contradiction. Next suppose A does not allocate M to j. Consider input instance $I_2 \subset I_1$ that contains only item $(j, M, 1+\delta)$. Then A cannot distinguish between I_1 and I_2 when considering item $(j, M, 1+\delta)$. Thus A will not allocate M to bidder j on input I_2, which contradicts Theorem 1.

Case 2: $(j, S, 1)$, $S \neq M$. Consider the decision made by A for this item. Suppose A does not allocate S to bidder j. Let $I_3 \subseteq I_1$ be the input instance consisting only of items $(j, T, 1)$ for all $T \supseteq S$; that is, player j has a single-minded valuation for set S. Since A cannot distinguish between I_1 and I_3 when considering item $(j, S, 1)$, it must be that A does not allocate S to bidder j on input I_3. Since A does not allocate any set T to player j other than set S (by assumption), it must not allocate anything to player j. Thus A obtains a social welfare of 0 when 1 was possible, contradicting the supposed approximation ratio of A.

Thus A must allocate S to bidder j on input I_3. Let $I_4 \subseteq I_1$ be the input instance consisting of items $(j, S, 1)$ and $(j, M, 1+\epsilon)$. Then A will allocate S to bidder j in instance I_4, but this contradicts Theorem 1 (which requires that A allocate M to bidder j).

We therefore arrive at a contradiction in all cases, as required. □

3.2 Bidders as Items

Roughly speaking, the lower bounds in Theorems 2 and 3 follow from a priority algorithm's inability to determine which of many different mutually-exclusive desires of an agent to consider first when constructing an allocation. One might guess that such difficulties can be overcome by presenting an algorithm with more information about an agent's valuation function at each step. To this end, we consider an alternative model of priority algorithms in which the agents themselves are the items, and the algorithm is given complete access to an agent's declared valuation function each round.

Under this model, $\mathcal{I}$ consists of all pairs (i, v_i), where $i \in N$ and $v_i \in V_i$. A valid input instance contains one item for each bidder. The decision to be made for item (i, v_i) is a set $S \subseteq M$ to assign to bidder i. The truthful greedy CA mechanism for single-minded bidders due to Lehmann et al. [40] falls within this model, as does its (non-truthful) generalization to complex bidders [40], the primal-dual algorithm of [12], and the (first) algorithm of [7] for multi-unit CAs. We now establish an inapproximation bound for truthful priority allocations that use bidders as items.

Theorem 4. *Suppose A is an incentive compatible priority algorithm for the (2-minded) CA problem that uses bidders as items. Than A cannot approximate the optimal social welfare by a factor of* $\frac{(1-\delta)k}{2}$ *for any* $\delta > 0$.

Proof. Choose $\delta > 0$ and suppose for contradiction that A is an incentive compatible adaptive priority algorithm that achieves an approximation ratio of $k(1-\delta)/2$. Recall that an item is a tuple (i, v_i), where $1 \leq i \leq n$ is a bidder and $v_i : 2^M \to \mathbb{R}$ is a valuation function.

We will construct a set of input instances for which A is forced to make a particular allocation, due to incentive compatibility. We define two sets of valuation functions, $\{g_1, \ldots, g_k\}$ and $\{f_1, \ldots, f_k\}$, that will be used in these input instances. The functions $g_1, \ldots, g_k$ are straightforward: for each $1 \leq i \leq k$, define valuation function g_i by

$$g_i(S) = \begin{cases} 1 & \text{if } a_i \in S \\ 0 & \text{otherwise.} \end{cases}$$

Then g_i is a single-minded valuation function, where the desired set is $\{a_i\}$ with value 1.

The definition of valuation functions $f_1, \ldots, f_k$ is more involved. Fix $i \in N$ and define $V'_{-i} := \{g_1, \ldots, g_k\}^{n-1}$. Consider an instance $\mathbf{d}$ of the combinatorial auction problem in which $\mathbf{d}_{-i} \in V'_{-i}$. That is, each bidder $j \neq i$ is single-minded, and desires a singleton with value 1. By the critical price property, there is a critical price $p_i(M, \mathbf{d}_{-i})$ for set M given this $\mathbf{d}_{-i}$.

Lemma 4. $p_i(M, \mathbf{d}_{-i}) \leq kn$.

Proof. Suppose otherwise that $p_i(M, \mathbf{d}_{-i}) > kn$. Suppose further that bidder i is single-minded with desired set M, and with $d_i(M) = kn$. Then $d_i(M) - p_i(M, \mathbf{d}_{-i}) < 0 = d_i(\emptyset) - p_i(\emptyset, \mathbf{d}_{-i})$. Therefore, by the critical pricing property, A cannot allocate M to bidder i, and hence bidder i obtains a value of 0. Now consider the social welfare obtained by A: it can be at most $n-1$, since bidder i obtains a welfare of 0 and each other bidder has value at most 1 for any set. The optimal social welfare is kn, obtained by allocating M to bidder i. Hence A obtains an approximation ratio of $\frac{kn}{n-1} > \frac{k(1-\delta)}{2}$ for this input instance, which is a contradiction. This completes the proof of Lemma 4.

We are now ready to define the valuations $f_1, \ldots, f_k$. They are based on values $x, y \in \mathbb{R}$. Define $x \in \mathbb{R}$ as follows:

$$x := 1 + \max_{i \in N} \max_{\mathbf{v}_{-i} \in V'_{-i}} \{p_i(M, \mathbf{v}_{-i})\}.$$

That is, x is a value greater than the maximum of the critical price for M for bidder i, over all choices of i and possible desires of singletons with value 1 by other bidders. Set $y := x\delta^{-1}$.

For each $1 \leq i \leq k$, define valuation function f_i as

$$f_i(S) = \begin{cases} y & \text{if } \{a_i\} \subseteq S \subset M \\ y + x & \text{if } S = M \\ 0 & \text{otherwise.} \end{cases}$$

Then $f_i(S)$ is a 2-minded valuation function. We now consider the following subset $\mathcal{I}' \subseteq \mathcal{I}$ of possible input items: $\mathcal{I}'$ contains all bidder-valuation pairs of the form (i, v_i) where $1 \leq i \leq n$ and $v_i = f_j$ or $v_i = g_j$ for some $1 \leq j \leq k$. Note that $\mathcal{I}'$ is not a valid input instance; we think of $\mathcal{I}'$ simply as a subset of $\mathcal{I}$.

The following claim exploits the incentive compatibility of A.

Lemma 5. *Suppose $I = \{(1, d_1), \ldots, (n, d_n)\}$ is a valid input instance, in which there exists $i \in N$ such that $d_i \in \{f_1, \ldots, f_k\}$, and for all $j \neq i$, $d_j \in \{g_1, \ldots, g_k\}$. Then on input I, A must allocate M to bidder i and $\emptyset$ to all other bidders.*

Proof. For this input instance we have that $\mathbf{d}_{-i} \in V'_{-i}$. Then $x > p_i(M, \mathbf{v}_{-i})$ from the definition of x. But now, from the definition of f_i,

$$d_i(M) - p_i(M, \mathbf{v}_{-i}) > (y + x) - x = y \geq d_i(S) \geq d_i(S) - p_i(S, \mathbf{d}_{-i})$$

for all $S \neq M$. Therefore, by the critical pricing property (Theorem 1), A must allocate M to bidder i, completing the proof of Lemma 5.

Our next step is to construct an input instance $I \subseteq \mathcal{I}'$ on which A obtains a poor approximation ratio. To do this we will rely on the following claim which will be proven by induction on i.

Lemma 6. *There exists a labelling of bidders and objects such that the following is true for all $0 \leq i < k/2$. Define $I_i := \{(j, g_j) \mid j \leq i\}$. Then for any valid input instance I such that $I_i \subseteq I \subseteq \mathcal{I}'$, A will consider all the items in I_i before all other items in I, and will choose to assign $\emptyset$ for each of the items in I_i.*

Proof. We proceed by induction on i. The base case holds by taking $I_0 = \emptyset$. For general $i \geq 1$, suppose the claim is true for $i - 1$. Then $I_{i-1} = \{(1, g_1), \ldots, (i - 1, g_{i-1})\}$. Define $\mathcal{I}_i \subseteq \mathcal{I}'$ as follows:

$$\mathcal{I}_i := \{(j, v_j) \mid (j, v_j) \in \mathcal{I}', j \geq i\}.$$

That is, $\mathcal{I}_i$ contains items of $\mathcal{I}'$ corresponding to bidders that are not present in I_{i-1}. Then note that if I is a valid input instance such that $I_{i-1} \subseteq I$, then $I \subseteq I_{i-1} \cup \mathcal{I}_i$.

Consider the execution of A on any valid input instance $I \subseteq I_{i-1} \cup \mathcal{I}_{i-1}$. The algorithm will first consider the items of I_{i-1} and allocate $\emptyset$ to each bidder $1, \ldots, i-1$ (by assumption). Once this is done, the algorithm will choose an ordering $\mathcal{T}$ over $\mathcal{I}_i$ and examine the next item in I according to $\mathcal{T}$.

Some item $(j, v_j) \in \mathcal{I}_i$ must come first under this ordering $\mathcal{T}$. Without loss of generality (by relabeling indices) this item is (i, f_i) or (i, g_i). We consider these two cases separately.

Case 1: The first item is (i, f_i). In this case we will choose I so that $(i, f_i) \in I$. Then A must consider this item next when processing input instance I, and A must assign some set S to bidder i. If $S = M$, then we will choose I to contain (j, f_j) for all $j > i$; note that $I \subseteq I_{i-1} \cup \mathcal{I}_{i-1}$ as required. Since A allocated M to bidder i, it obtains a social welfare of $x+y$ on input I. However, the optimum welfare is at least $(k-i+1)y$, since this can be attained by allocating $\{a_j\}$ to bidder j for all $i \leq j \leq k$. Thus the approximation ratio obtained by A is at least

$$\frac{(k-i+1)y}{x+y} > \frac{(k/2)y}{y(1+\delta)} > \frac{(1-\delta)k}{2},$$

a contradiction.

If, on the other hand, $S \neq M$, we choose I to contain (j, g_j) for all $j > i$. Then I satisfies the requirements of Lemma 5, so A must allocate M to bidder i. This is a contradiction. We conclude that this first case cannot occur.

Case 2: The item is (i, g_i). In this case we will choose I so that $(i, g_i) \in I$. As in the previous case, A must consider this item next in I, and assign some set S to bidder i. Suppose $S \neq \emptyset$. Then we will choose I to contain $(i+1, f_{i+1})$, and also (j, g_j) for all $j > i+1$. Note that then $I \in I_{i-1} \cup \mathcal{I}_{i-1}$ as required. Also, in this instance of a combinatorial auction, $v_{-(i+1)}$ contains only single-minded valuations for singletons with value 1. Thus, by the same argument used in Case 1, it must be that bidder $i+1$ is allocated M. However, this is not possible, since bidder i is assigned $S \neq \emptyset$. This is a contradiction. We conclude that in this case, bidder i must be assigned $\emptyset$.

This ends our case analysis. We conclude that item (i, g_i) must occur first in $\mathcal{I}_{i-1}$ in ordering $\mathcal{T}$, and furthermore if $(i, g_i) \in I$ then A will consider (i, g_i) next after processing the items in I_{i-1} and will assign $\emptyset$ to bidder i. We can therefore set $I_i = I_{i-1} \cup \{(i, g_i)\}$ to satisfy the requirements of the Lemma 6.

Now suppose $I_{k/2-1}$ is the set from Lemma 6 with $i = k/2 - 1$. Define input instance I by

$$I := I_{k/2-1} \cup \{(j, g_{k/2}) \mid k/2 \leq j \leq k\}.$$

Note that I is a valid input instance and $I_{k/2-1} \subseteq I \subseteq \mathcal{I}'$. Then by Lemma 6, algorithm A must assign $\emptyset$ to each of bidders $1, \ldots, k/2-1$. Therefore A can obtain a social welfare of at most 1, by assigning $\{a_{k/2}\}$ to some bidder $j \geq k/2$. However, the optimal social welfare is $k/2$, by assigning $\{a_i\}$ to bidder i for all

$1 \leq i \leq k/2$. Hence A obtains an approximation no better than $k/2$, which is a contradiction. This completes the proof of Theorem 4. □

4 Truthful Submodular Priority Auctions

Lehmann, Lehmann, and Nisan [39] proposed a class of greedy algorithms that is well-suited to auctions with submodular bidders; namely, objects are considered in any order and incrementally assigned to greedily maximize marginal utility. They showed that any ordering of the objects leads to a 2-approximation of social welfare, but not every ordering of objects leads to an incentive compatible algorithm. However, this does not preclude the possibility of obtaining truthfulness using some adaptive method of ordering the objects.

We consider a model of priority algorithms which uses the m objects as input items. In this model, an item will be represented by an object x, plus the value $v_i(x|S)$ for all $i \in N$ and $S \subseteq M$ (where $v_i(x|S) := v_i(S \cup \{x\}) - v_i(S)$ is the marginal utility of bidder i for item x, given set S). We note that the online greedy algorithm described above falls into this model. We show that no greedy priority algorithm in this model is incentive compatible.

Theorem 5. *Any greedy priority algorithm for the combinatorial auction problem that uses objects as items is not incentive compatible. This holds even if the bidders are assumed to be submodular.*

Proof. Suppose for contradiction that A is an incentive compatible truthful greedy priority algorithm. Consider an instance of the combinatorial auction with $M = \{a_1, a_2, a_3\}$. Suppose that bidder 1 declares the following valuation function: $v_1(S) = 9 + |S|$ for all $S \neq \emptyset$. It is easy to verify that this is indeed submodular. Then by Theorem 1 this valuation defines a critical price $p_2(S)$ for each subset $S \subseteq M$. Consider the critical prices for all subsets of size 2 and suppose without loss of generality that $\{a_1, a_2\}$ has the smallest. That is, $p_2(\{a_1, a_2\}) \leq p_2(\{a_2, a_3\})$ and $p_2(\{a_1, a_2\}) \leq p_2(\{a_1, a_3\})$.

We now define a valuation function v_2 to be declared by bidder 2. The motivation for v_2 is that items a_1 and a_2 will have lower values than a_3 when considered individually, but will have a large value relative to a_3 when taken together.

$$
\begin{aligned}
v_2(\{a_1\}) = v_2(\{a_2\}) &= 9 \\
v_2(\{a_3\}) &= 11 \\
v_2(\{a_1, a_2\}) &= 18 \\
v_2(\{a_1, a_3\}) = v_2(\{a_2, a_3\}) &= 17 \\
v_2(\{a_1, a_2, a_3\}) &= 18.
\end{aligned}
$$

It is easily verified that this valuation is submodular.

Given as input the valuations v_1 and v_2, algorithm A must consider each object in turn, and assign that object to the player who obtains the greatest marginal utility from it. The algorithm is free to choose the order in which the items are considered. However, regardless of the order, the only possible outcomes are that bidder 2 is allocated $\{a_1, a_3\}$ or bidder 2 is allocated $\{a_2, a_3\}$. This can be seen by examining each of the 6 possible orderings of items, or by noticing that the first item considered will go to bidder 2 if and only if it is a_3, that bidder 1 will never be allocated the second object considered, and that bidder 2 will never be allocated the third object considered.

We will assume without loss of generality that bidder 2 is allocated $\{a_1, a_3\}$. Then, by Theorem 1,

$$v_2(\{a_1, a_3\}) - p_2(\{a_1, a_3\}) \geq v_2(\{a_1, a_2\}) - p_2(\{a_1, a_2\}).$$

Since $v_2(\{a_1, a_3\}) = 17$ and $v_2(\{a_1, a_2\}) = 18$, this implies that

$$p_2(\{a_1, a_2\}) > p_2(\{a_1, a_3\})$$

which contradicts the minimality of $p_2(\{a_1, a_2\})$.

We have now proved the result for the case of exactly two bidders and three objects. The result follows in the desired generality by noticing that we may add additional players who value all sets at 0, and additional items for which no players have value, without affecting the above construction.

5 Future Work

The goal of algorithmic mechanism design is the construction of algorithms in situations where inputs are controlled by selfish agents. We considered this fundamental issue in the context of conceptually simple methods (independent of time bounds) rather than in the context of time constrained algorithms. Our results concerning priority algorithms (as a model for greedy mechanisms) is a natural beginning to a more general study of the power and limitations of conceptually simple mechanisms. Even though the priority framework represents a restricted (albeit natural) algorithmic approach, there are still many unresolved questions even for the most basic mechanism design questions. In particular, we believe that the results of Section 3 can be unified to show that the linear inapproximation bound holds for all priority algorithms for s-CA problems. The power of greedy algorithms for unit-demand auctions (s-CAs with $s = 1$) is also not understood. While there are polynomial time (i.e. edge weighted bipartite matching) algorithms, it is not difficult to show that optimality cannot be achieved by priority algorithms. But is it possible to obtain a sublinear truthful approximation bound for 1-CAs with greedy methods?

An obvious direction of future work is to widen the scope of a systematic search for truthful approximation algorithms; priority algorithms can be extended in many ways. Perhaps the most immediate extension is to consider randomized priority algorithms (as in [1]) for the CA problem. The currently best known randomized truthful (and truthful in expectation) mechanisms with sublinear approximation ratios are not greedy algorithms. One might also consider priority algorithms with a more esoteric input model, such as a hybrid of the sets as items and bidders as items models. Priority algorithms can be extended to allow revocable acceptances [33] whereby a priority algorithm may "de-allocate" sets or objects that had been previously allocated to make a subsequent allocation feasible. Somewhat related is the priority stack model [10] (as a formalization of local ratio/primal dual algorithms with reverse delete [6]) where items (e.g. bidders or bids) initially accepted are placed in a stack and then the stack is popped to ensure feasibility. This is similar to algorithms that allow a priority allocation algorithm to be followed by some simple "cleanup" stage [35]. Another possibility is to consider allocations that are comprised of taking the best of two (or more) priority algorithms. A special case that has been used in the design of efficient truthful combinatorial auction mechanisms [7, 12, 45] is to optimize between a priority allocation and the naïve allocation that gives all objects to one bidder. Finally, one could study more general models for algorithms that implement integrality gaps in LP formulations of packing problems; it would be of particular interest if a deterministic truthful k-approximate mechanism could be constructed from an arbitrary packing LP with integrality gap k, essentially derandomizing the construction of Lavi and Swamy [38].

The results in this paper have thus far been restricted to combinatorial auctions but *the basic question* being asked applies to other algorithmic mechanism design problems such as machine scheduling or more general integer programming problems. Namely, when can a conceptually simple approximation to the underlying combinatorial optimization problem be converted into an incentive compatible mechanism that achieves (nearly) the same approximation? For example, one might consider the power of truthful priority mechanisms for approximating unrelated machines scheduling, or for more general integer packing problems.

Acknowledgements

We thank S. Dobzinski, R. Gonen, and S. Micali for helpful discussions.

References

1. S. Angelopoulos and A. Borodin. Randomized priority algorithms. *Theoretical Computer Science*, 411:2542–2558, 2010.
2. A. Archer, C. Papadimitriou, K. Talwar, and E. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *Proc. 14th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2003.
3. L. M. Ausubel and P.R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1), 2002.
4. Y. Azar, I. Gamzu, and S. Gutner. Truthful unsplittable flow for large capacity networks. In *Proc. 19th ACM Symp. on Parallel Algorithms and Architectures*, 2007.
5. M. Babaioff and L. Blumrosen. Computationally-feasible truthful auctions for convex bundles. In *Proc. 7th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2004.

6. A. Bar-Noy, R. Bar-Yehuda, A. Freund, J. Naor, and B. Schieber. A unified approach to approximating resource allocation and scheduling. *Journal of the ACM*, 48:1069–1090, 2001.
7. Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *Proc. 9th Conf. on Theoretical Aspects of Rationality and Knowledge*, 2003.
8. X. Bei and Z. Huang. Towards optimal bayesian algorithmic mechanism design. In *Proc. 22nd ACM Symp. on Discrete Algorithms*, 2011.
9. K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proc. 22nd ACM Symp. on Discrete Algorithms*, 2011.
10. A. Borodin, D. Cashman, and A. Magen. How well can primal-dual and local-ratio algorithms perform? In *Proc. 32nd Intl. Colloq. on Automata, Languages and Programming*, 2005.
11. A. Borodin, M. Nielsen, and C. Rackoff. (Incremental) priority algorithms. *Algorithmica*, 37:295–326, 2003.
12. P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanism design. In *Proc. 36th ACM Symp. on Theory of Computing*, 2005.
13. D. Buchfuhrer, S. Dughmi, H. Fu, R. Kleinberg, E. Mossel, C. Papadimitriou, M. Schapira, Y. Singer, and C. Umans. Inapproximability for vcg-based combinatorial auctions. In *Proc. 21st ACM Symp. on Discrete Algorithms*, 2010.
14. C. Chekuri and I. Gamzu. Truthful mechanisms via greedy iterative packing. In *Proc. 12th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2009.
15. G. Christodoulou, A. Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *Proc. 35th Intl. Colloq. on Automata, Languages and Programming*, pages 820–832, 2008.
16. P. Cramton. Spectrum auctions. In M. Cave, S. Majumdar, and I. Vogelsang, editors, *Handbook of Telecommunications Economics*, chapter 14, pages 605–639. Elsevier Science B. V., 2002.
17. P. Cramton, Y. Shoham, and R. Steinberg (eds.). *Combinatorial Auctions.* MIT Press, Cambridge, MA, 2005.
18. P. Dhangwatnotai, S. Dobzinski, S. Dughmi, and T. Roughgarden. Truthful approximation schemes for single-parameter agents. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, 2008.
19. S. Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Proc. 11th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2008.
20. S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *STOC11*, 2011. to appear.
21. S. Dobzinski and S. Dughmi. On the power of randomization in algorithmic mechanism design. In *Proc. 50th IEEE Symp. on Foundations of Computer Science*, 2009.
22. S. Dobzinski and N. Nisan. Limitations of vcg-based mechanisms. In *Proc. 38th ACM Symp. on Theory of Computing*, 2007.
23. S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proc. 37th ACM Symp. on Theory of Computing*, 2006.
24. S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Mathematics of Operations Research*, 35(1):1–13, 2010.

25. S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proc. 17th ACM Symp. on Discrete Algorithms*, 2006.
26. S. Dobzinski and M. Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *Proc. 10th ACM Conf. on Electronic Commerce*, 2008.
27. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In *Stanford Graduate School of Business Research Paper No. 1917*, 2005.
28. R. Gonen and D. Lehmann. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proc. 2nd ACM Conf. on Electronic Commerce*, 2000.
29. J. Hartline, R. Kleinberg, and A. Malekian. Bayesian incentive compatibility and matchings. In *Proc. 22nd ACM Symp. on Discrete Algorithms*, 2011.
30. J. Hartline and B. Lucier. Bayesian algorithmic mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing*, 2010.
31. J. Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. *Electronic Colloquium on Computational Complexity*, 38, 1997.
32. R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatiorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.
33. S. Horn. One-pass algorithms with revocable acceptances for job interval selection. University of Toronto MSC thesis, 2004.
34. S. Khot, R. Lipton, E. Markakis, and A. Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. In *Workshop on Internet and Network Economics (WINE)*, 2005.
35. P. Krysta. Greedy approximation via duality for packing, combinatorial auctions and routing. In *Proc. 30th Intl. Symp. on Mathematical Foundations of Computer Science*, 2005.
36. R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proc. 44th IEEE Symp. on Foundations of Computer Science*, 2003.
37. R. Lavi and N. Nisan. Online ascending auctions for gradually expiring goods. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.
38. R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, 2005.
39. B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proc. 3rd ACM Conf. on Electronic Commerce*, 2001.
40. D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce*, pages 96–102. ACM Press, 1999.
41. R. Paes Leme and E. Tardos. Sponsored search equilibria for conservative bidders. In *Fifth Workshop on Ad Auctions*, 2009.
42. B. Lucier. Beyond equilibria: Mechanisms for repeated combinatorial auctions. In *Proc. 1st Symp. on Innovations in Computer Science*, 2010.
43. B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proc. 21st ACM Symp. on Discrete Algorithms*, 2010.
44. M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. In *Proc. 8th ACM Conf. on Electronic Commerce*, 2006.

45. A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64:612–631, 2008.
46. N. Nisan. Bidding and allocation in combinatorial auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce*, 2000.
47. C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, 2008.
48. S. J. Rassenti, V. L. Smith, and R. L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13:402–417, 1982.
49. D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proc. 37th ACM Symp. on Theory of Computing*, 2006.