On the Limitations of Greedy Mechanism Design for Truthful Combinatorial Auctions

ALLAN BORODIN, University of Toronto
BRENDAN LUCIER, Microsoft Research

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 polynomial time deterministic truthful mechanism is known to attain an approximation ratio better than $O(m/\sqrt{\log m})$. We view this not only as an extreme gap between the power of greedy auctions and truthful greedy auctions, but also as exemplifying the gap between the known power of truthful and non-truthful polynomial time deterministic algorithms. We associate the notion of greediness with a broad class of algorithms, known as priority algorithms, which encapsulate 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 $s \geq 2$. Our inapproximations are independent of any time constraints on the mechanism and are purely a consequence of the greedy-type restriction. We conclude that any truthful combinatorial auction mechanism with a non-trivial approximation factor must fall outside the scope of many natural auction methods.

CCS Concepts: • **Theory of computation → Algorithmic game theory and mechanism design**; **Algorithmic mechanism design**; *Theory and algorithms for application domains;*

Additional Key Words and Phrases: Greedy algorithms, combinatorial auctions, truthfulness

ACM Reference Format:
Allan Borodin and Brendan Lucier. 2016. On the limitations of greedy mechanism design for truthful combinatorial auctions. ACM Trans. Econ. Comput. 5, 1, Article 2 (October 2016), 23 pages.
DOI: http://dx.doi.org/10.1145/2956585

1. INTRODUCTION

As introduced in the seminal paper of Nisan and Ronen [2001], 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 Vickrey-Clark-Groves (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

Authors' addresses: A. Borodin, Department of Computer Science, University of Toronto, 10 Kings College Rd., Toronto, Ontario, Canada M5S3G4; email: bor@cs.toronto.edu; B. Lucier, Microsoft Research, New England, One Memorial Drive, Cambridge, Mass. 02142, USA; email: brlucier@microsoft.com.

DOI: http://dx.doi.org/10.1145/2956585

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 different 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 and substitutes; for example, an agent's complementary 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, and two left shoes could certainly be worth less than twice the value of one left shoe. Combinatorial auctions have arisen in practice for the sale of airport landing schedules [Rassenti et al. 1982], FCC spectrum auctions [Cramton 2002], and others; see Cramton et al. [2005] for an overview. The CA problem, also known as the welfare maximization problem, is to determine, given the agents' valuation functions (either explicitly 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 [Lehmann et al. 2002], and this is the best possible under standard complexity assumptions [Hastad 1999; Zuckerman 2007]. However, no deterministic truthful mechanism is known to obtain an approximation ratio better than $O(\frac{m}{\sqrt{\log m}})$ for the general problem [Holzman et al. 2004]. This is true even for the natural and broad class of valuations where each bidder is interested only in sets of size at most some constant $s \geq 2$ (the s-CA problem)[1], a natural and broad class of valuations for which the standard greedy algorithm obtains an $s + 1$ approximation. Understanding when these gaps are essential for the general CA problem, and for the CA problem when restricted to natural classes of agent valuation functions, is a central open question that has received significant attention over the past decade [Dobzinski and Sundararajan 2008; Lavi et al. 2003; Lehmann et al. 2002; Mu'alem and Nisan 2008; Papadimitriou et al. 2008]. As we discuss further in Section 1.2, there have been some recent impossibility results in this regard for important classes of valuations, including submodular valuations [Dobzinski 2011; Dobzinski and Vondrák 2012]. Most recently, in a significant development, Daniely et al. [2015] have shown that for every $\epsilon > 0$, there exists a class of valuations that can be m^{ϵ} (non-truthfully) approximated but for which no truthful mechanism can provide an $(1 - \epsilon)$ approximation unless NP is a subset of the non-uniform analogue of P (i.e., $P/poly$). Nevertheless, the fundamental problem still remains open for many valuation classes of interest such as s-CAs.

The hope, when studying the combinatorial auction problem, would be to find a natural and truthful approximation mechanism. This qualifier "natural" is highly subjective, but nevertheless important; examples that come to mind are the well-known socially optimal Vickrey auction for a single object, or item pricing for revenue maximization in the full information setting [Guruswami et al. 2005] and sequential posted pricing in the Bayesian setting [Chawla et al. 2010]. 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 [Ausubel and Milgrom 2002]. As Syrgkanis and Tardos [Syrgkanis and Tardos 2013] state, "the Internet environment allows for running millions of auctions, which necessitates the use of very simple and intuitive auction

[1]For specific interest in such small bundle CAs, see, for example, Kesselheim et al. [2013].

schemes." We are therefore motivated to ask the following loosely-defined question: can any "natural" auction, which 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? To this end, we will study the approximations that can be achieved by truthful greedy algorithms.

Since the standard greedy algorithm for the combinatorial auction problem is not truthful [Lehmann et al. 2002], 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. [2003] for the multi-unit combinatorial auction problem is a primal-dual algorithm that proceeds by iteratively constructing a price vector. Their algorithm 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 [Cramton et al. 2005]. One might hope that such methodologies, extended further, could lead to similar greedy-like truthful algorithms for common types of combinatorial auction problems.

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 $O(\frac{m}{\sqrt{\log(m)}})$ or even better than $O(m)$? Our specific interest in greedy algorithms is motivated threefold. First, most known examples of truthful, non-MIR (maximal in range) algorithms for combinatorial auction problems apply greedy methods [Azar et al. 2010; Bartal et al. 2003; Briest et al. 2011; Chekuri and Gamzu 2009; Dütting et al. 2014; Krysta 2005; Lehmann et al. 2006, 2002; Mu'alem and Nisan 2008]; 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 [Edelman et al. 2005]). 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, simplicity, perceived fairness, and computational efficiency.

Our article is organized as follows. The remaining subsections of Section 1 informally state our results and provide further discussion of relevant work. Section 2 provides the necessary definitions as well as a critical-price characterization of truthful mechanisms due to Bartal et al. [2003] and a variation thereof. Section 3 contains our main results, establishing the limitation of greedy algorithms (formalized as priority algorithms) for the general CA and s-CA problems. Section 4 discusses greedy auctions for CAs with submodular valuations, and we conclude with a discussion of some open problems in Section 5.

1.1. Our Results

As stated, our focus is on greedy and "greedy-like" algorithms, and to that end, we need to formulate precise definitions for the class of algorithms we will consider. We use the term "greedy algorithm" to refer to any of a large class of algorithms known as *priority algorithms* [Borodin et al. 2003]. The class of priority algorithms captures a general notion of "myopic algorithm" behavior.[2] We review the priority framework

[2]The term *myopic algorithm* precedes the more commonly used terminology of greedy algorithms. Specifically, myopic algorithms do not presume that the decisions that need to be made (e.g., whether or not to accept

in Section 2.3. In Section 3, we adapt and apply the priority model to the problem of direct revelation mechanisms for combinatorial auctions. Informally speaking, in our context, a priority algorithm proceeds by ranking bids according to some (possibly adaptive) quality score; winning a certain bundle in some iteration requires that one's bid does not conflict with any previously allocated items and that this bid be ranked higher than other competing bids.

Our main result demonstrates that if a truthful auction for a CA falls within the priority framework, then it cannot, in general, perform much better than the naïve algorithm that allocates all objects to a single bidder. That is, its approximation ratio will be $\Omega(\min\{n, m\})$, where n is the number of agents and m is the number of objects in the auction. We note that this inapproximation result does not depend on computational or communication complexity assumptions, and is incomparable with inapproximations based on computational hardness. The gap described in our result is extreme for s-CAs: for $s = 2$, the standard (but non-truthful) greedy algorithm is a 3-approximation for the multi-minded s-CA problem,[3] but no truthful greedy algorithm can obtain a sublinear (in n and m) approximation bound.

We also consider the combinatorial auction problem for submodular bidders (SMCA), a 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 Lehmann et al. [2006] that any such algorithm 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.[4]

1.2. Related Work

It is known that there 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 in this regard was first proven for the related combinatorial public project problem [Papadimitriou et al. 2008]. They show that there is a large asymptotic gap separating approximation by deterministic algorithms and by deterministic truthful mechanisms in general allocation problems.

Significant work has focused on particular restricted cases of the combinatorial auction problem, such as submodular auctions [Dobzinski and Schapira 2006; Dobzinski et al. 2010; Khot et al. 2008; Lehmann et al. 2006; Dobzinski 2011], and on alternative solution concepts such as randomized notions of truthfulness [Dhangwatnotai et al. 2011; Lavi and Swamy 2011; Dobzinski and Dughmi 2013; Archer et al. 2003; Dobzinski et al. 2012; Dobzinski 2007], truthfulness in Bayesian settings [Hartline and Lucier 2010; Hartline et al. 2011; Bei and Huang 2011], and performance at (non-truthful) equilibrium [Lucier and Borodin 2010; Lucier 2010; Syrgkanis and Tardos 2013; Caragiannis et al. 2015]. The relatively recent results of Dobzinski [2011] and Dobzinski and Vondrák [2012] have established a large gap between non-truthful and

a bid) will be made greedily (e.g., that a bid, when considered, will be accepted if it does not conflict with previous decisions).

[3]The standard greedy algorithm greedily selects non-conflicting bids in order of non-increasing value. The underlying allocation problem is the s-set packing (respectively, $s+1$-set packing) problem for single-minded (respectively, multi-minded) bidders. This problem is a special case of the maximum independent set problem in $s+1$ (respectively, $s+2$) claw-free graphs for which the standard greedy algorithm provides an s-approximation (respectively, an $s+1$-approximation). See Chandra and Halldórsson [2001].

[4]The degree of freedom in this class of algorithms is the order in which the objects are considered. We note that our conference paper [Borodin and Lucier 2010] preceded the very strong computational-complexity-based inapproximation results for truthful submodular CAs given in Dobzinski [2011] and Dobzinski and Vondrák [2012], discussed in more detail in Section 1.2.

truthful mechanisms for the bidders with submodular valuations. Namely, assuming $NP = RP$ (respectively, NP is not contained in non-uniform polynomial time), truthful deterministic algorithms cannot achieve an $m^{\frac{1}{2}-\epsilon}$ approximation for any $\epsilon > 0$ (respectively, approximation n^{γ} for some $\gamma > 0$). In contrast, without the game-theoretic truthfulness requirement, there is a polynomial time $\frac{e}{e-1}$ approximation algorithm and a 2-approximation greedy algorithm for the CA problem when agents all have submodular valuation functions. Most recently, in a significant development, Daniely et al. [2015] have shown that for every $\epsilon > 0$, there exists a class of valuations that can be m^{ϵ} (non-truthfully) approximated but for which no truthful mechanism can provide an $m^{1-\epsilon}$ approximation unless NP is a subset of the non-uniform analogue of P (i.e., $P/poly$). Nevertheless, beyond closing the gap between the $m^{1-\epsilon}$ inapproximation and the $O(\frac{m}{\sqrt{\log m}})$ truthful deterministic mechanism, the problem for specific classes of CAs (such as S-CAs) still 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 [Ausubel and Milgrom 2002]. 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.

There have been many developments in the restricted case of CAs with single-minded bidders. Following the Lehmann et al. [2002] truthful greedy mechanism for single-minded CAs, Mu'alem and Nisan [2008] 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 [Babaioff and Blumrosen 2008; Briest et al. 2011] and additional construction techniques, such as the iterative greedy packing method due to Chekuri and Gamzu [2009].

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 $O(m/\sqrt{\log m})$ [Holzman et al. 2004]. For the special case of multi-unit CAs, when there are $B \geq 3$ copies of each object, Bartal et al. [2003] give a greedy algorithm that obtains an $O(Bm^{\frac{1}{B-2}})$ approximation. Lavi and Swamy [2011] 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. These verifiers take the form of greedy algorithms, which play a prominent role in the final mechanisms. Dobzinski et al. [2010] construct a universally truthful randomized $O(\sqrt{m})$-approximate mechanism for the CA problem via sampling.

Prior to the inapproximation results of Dobzinski [2011] and Dobzinski and Vondrák [2012], a substantial line of research gave lower bounds on the approximating power of somewhat restricted classes of deterministic truthful algorithms for CAs. Lavi et al. [2003] 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 two cannot satisfy the natural property of being *stable*[5] [Dobzinski and Sundararajan 2008]. 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 [Dobzinski and Nisan 2011]. 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) [Buchfuhrer et al. 2010]. These lower bounds are incomparable to our own, as priority algorithms need not be MIR, stable, unanimity-respecting, or satisfy IIA.[6]

There has been extensive work studying the power of truthful mechanisms for restricted forms of combinatorial auctions, such as submodular auctions [Dobzinski and Schapira 2006; Dobzinski et al. 2010; Khot et al. 2008; Lehmann et al. 2006; Dobzinski 2011; Dobzinski et al. 2012]. Of particular relevance to our work is the recent work of Dobzinski [Dobzinski 2011] 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 [Lehmann et al. 2006].

Another line of work gives lower bounds for greedy algorithms without truthfulness restrictions. Gonen and Lehmann [2000] showed that no algorithm that greedily accepts bids for sets can guarantee an approximation better than $\sqrt{m}$ for the general CA problem. More generally, Krysta [2005] 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 more general class of priority algorithms but restrict them to be incentive compatible.

The class of priority algorithms substantially generalizes 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 [Lavi and Nisan 2015; Mahdian and Saberi 2006]. 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 adversarially. For example, the $O(\sqrt{m})$-approximate greedy algorithm for the CA problem, due to Lehmann et al. [2002], requires a specific ordering and cannot be achieved by an online algorithm. The multi-unit auction, due to Bartal et al. [2003], can be applied to online settings, though this requires that the mechanism have *a priori* bounds on the possible valuations, and the resulting approximation depends on the ratio between the minimum and maximum possible valuations.

In contrast to the negative results of this article, (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 [Christodoulou et al. 2008], and a similar auction method obtains a 2-approximation at Bayes-Nash equilibrium for subadditive bidders [Bhawalkar and Roughgarden 2011; Feldman et al. 2013]. The greedy GSP auction for internet advertising has also been shown to obtain a constant

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

[6]The notion of IIA has been associated with priority algorithms, but in a different context than in Lavi et al. [2003]. 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.

approximation to the optimal welfare at any Bayes-Nash equilibrium [Caragiannis et al. 2015]. It is also known that c-approximate greedy algorithms for combinatorial allocation problems can be converted into mechanisms whose Bayes-Nash equilibria yield $c(1+o(1))$ approximations [Lucier and Borodin 2010; Syrgkanis and Tardos 2013].

2. DEFINITIONS AND PRELIMINARY RESULTS

2.1. Combinatorial Auctions

A *combinatorial auction* consists of n bidders and a set M of m objects.[7] For notational convenience, we will let $k = \min\{m, n\}$. Each bidder i has a value for each subset of objects $S \subseteq M$, described by a valuation function $v_i : 2^M \rightarrow \mathbb{R}$, which we call the *type* of agent i. We assume each v_i is monotone and normalized, so that $v_i(S) \leq v_i(T)$ for all $S \subseteq T$ and $v_i(\emptyset) = 0$. 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 *ℓ-minded* if it is the maximum of ℓ single-minded functions. That is, there exist ℓ sets $S_1, \ldots , S_\ell$ 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 $S, T \subseteq M$. In the combinatorial auction problem, we are interested in *feasible allocations* which satisfy the property that no object is allocated to more than one bidder and each bidder is allocated at most one set. We sometimes refer to a feasible allocation as a *valid* allocation. An s-CA is one in which the valuation function of every agent satisfies $v(T) = \max_{S:|S| \leq s} val(S)$; that is, the agents only value sets of size at most s (and by monotonicty of valuations, all extension of such sets).

A *direct revelation mechanism* (or just *mechanism*) $\mathcal{M} = (A, P)$ consists of an *allocation algorithm* A 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 $\mathbf{d}$, $A(\mathbf{d})$ returns an allocation of objects to bidders, and $P(\mathbf{d})$ returns the payment extracted from each agent. For each agent i, we write $A_i(\mathbf{d})$ and $P_i(\mathbf{d})$ for the set given to and payment extracted from i. We will restrict payments to be non-negative.

The *social welfare* obtained by A on declaration $\mathbf{d}$ given truthful valuation $\mathbf{t}$ is $SW_A(\mathbf{d}, \mathbf{t}) = \sum_{i \in N} t_i(A_i(\mathbf{d}))$. 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 A is a c-approximation if $SW_A(\mathbf{t}, \mathbf{t}) \geq \frac{1}{c} SW_{opt}$ 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 $u_i(\mathbf{d}, \mathbf{t}) = t_i(A_i(\mathbf{d})) - P_i(\mathbf{d})$. When there is no confusion, we will often drop the dependence on $\mathbf{t}$ and simply write $u_i(\mathbf{d})$. Mechanism $\mathcal{M}$ is said to be *incentive compatible* (or *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 A is truthful if there exists a payment function P such that the mechanism (A, P) is truthful. We also say that $\mathcal{M}$ is *individually rational* if

[7] We are following the terminology for priority algorithms and reserve the word "item" to mean a basic unit of input. We then use "objects" to denote the goods being sold.

$u_i(t_i, \mathbf{d}_{-i}) \geq 0$ for all t_i and all $\mathbf{d}_{-i}$. That is, a rational agent does not strictly prefer to exclude himself from the mechanism.

2.2. Minimal Prices

Following Bartal, Gonen, and Nisan [2003], we will use a multi-parameter analogy of *critical prices* in order to characterize truthful mechanisms. Namely, an allocation algorithm A defines *minimal prices*, $p_i(S, \mathbf{d}_{-i})$, for any agent i and set S as follows. We define $p_i(\emptyset, \mathbf{d}_{-i}) = 0$ for all $\mathbf{d}_{-i}$. For $S \neq \emptyset$, the price $p_i(S, \mathbf{d}_{-i})$ is the minimum amount that agent i could bid on a set that contains S and win it, given that the other agents bid according to $\mathbf{d}_{-i}$. That is,

$$p_i(S, \mathbf{d}_{-i}) = \inf_{d_i : A_i(d_i, \mathbf{d}_{-i}) \supseteq S} \{d_i(A_i(d_i, \mathbf{d}_{-i}))\}.$$

Note that $p_i(S, \mathbf{d}_{-i})$ need not be finite. If $p_i(S, \mathbf{d}_{-i}) = \infty$, this indicates that the mechanism will simply not allocate S to bidder i for any reported valuation d_i, given that the other agents declare according to $\mathbf{d}_{-i}$. Moreover, it follows from the definition that, for all i, all $\mathbf{d}_{-i}$, and all $S \subseteq T$, $p_i(S, \mathbf{d}_{-i}) \leq p_i(T, \mathbf{d}_{-i})$.

The following characterization of truthful mechanisms for combinatorial auctions is due to Bartal et al. [2003]. As our notation and definitions are slightly different, we restate it here for completeness. Roughly speaking, a mechanism is truthful if the payment of an agent i is determined only by the declarations of the other agents and the allocation to agent i, and moreover, agent i is assigned the utility-maximizing allocation given these payments.

THEOREM 2.1 ([BARTAL ET AL. 2003]). *A mechanism $\mathcal{M} = (A, P)$ is truthful if and only if, for every bidder i, there is a function $\pi_i : 2^M \times V_{-i} \rightarrow \boldsymbol{R} \cup \{\infty\}$ such that, for all* $\mathbf{d}$,

(1) $P_i(\mathbf{d}) = \pi_i(A_i(\mathbf{d}), \mathbf{d}_{-i})$ *and*
(2) $A_i(\mathbf{d}) \in \text{argmax}_S\{d_i(S) - \pi_i(S, \mathbf{d}_{-i})\}$.

One can strengthen Theorem 2.1 if one assumes some structure on the class of valuations. Given a mechanism $\mathcal{M} = (A, P)$ and space of valuations V, we will write $R_i(A, V) = \{A_i(\mathbf{v}) : \mathbf{v} \in V\}$ for the range of outcomes that can be allocated to agent i. We will show that if the valuation space includes all single-minded declarations for outcomes in the range of the mechanism, and if we also require individual rationality, then, in fact, the function π_i from Theorem 2.1 must be the minimal price function p_i. For example, this condition would be satisfied by any mechanism for the general CA problem, as well as the s-CA problem (where each bidder can be allocated at most s objects and has arbitrary valuation over sets of size at most s).

THEOREM 2.2. *Fix a mechanism $\mathcal{M} = (A, P)$ and valuation class V such that, for each i, V_i includes all single-minded declarations for sets in $R_i(A, V)$. Then mechanism $\mathcal{M}$ is truthful and individually rational if and only if*

(1) $P_i(\mathbf{d}) = p_i(A_i(\mathbf{d}), \mathbf{d}_{-i})$ *and*
(2) $A_i(\mathbf{d}) \in \text{argmax}_S\{d_i(S) - p_i(S, \mathbf{d}_{-i})\}$.

PROOF. That the conditions are sufficient for truthfulness follows immediately from Theorem 2.1. The conditions are also sufficient for individual rationality because each agent's valuation class includes the zero function, and from the definition of p_i, we have $p_i(A_i(0, \mathbf{d}_{-i}), \mathbf{d}_{-i}) = 0$ for every $\mathbf{d}_{-i}$. The second condition of the theorem therefore implies that an agent's utility cannot be negative.

To prove the conditions are necessary, we first show that $P_i(\mathbf{d}) = p_i(A_i(\mathbf{d}), \mathbf{d}_{-i})$. Suppose $P_i(\mathbf{d}) > p_i(A_i(\mathbf{d}), \mathbf{d}_{-i})$ for some $\mathbf{d}$. Then, from the definition of minimal prices, there exists some d_i' such that $A_i(d_i', \mathbf{d}_{-i}) = A_i(\mathbf{d})$ and $d_i'(A_i(\mathbf{d})) < P_i(d)$. But then, if player i has type d_i', we have $u_i(d_i', \mathbf{d}_{-i}) < 0$, violating individual rationality.

Next, suppose $P_i(\mathbf{d}) < p_i(A_i(\mathbf{d}), \mathbf{d}_{-i})$ for some $\mathbf{d}$. Consider valuation d_i^* defined to be a single-minded declaration for set $A_i(\mathbf{d})$ at some value strictly between $P_i(\mathbf{d})$ and $p_i(A_i(\mathbf{d}), \mathbf{d}_{-i})$. Suppose that agent i has type d_i^*. From the definition of minimal prices, it must be that $A_i(d_i^*, \mathbf{d}_{-i}) \not\supseteq A_i(\mathbf{d})$, and hence, $u_i(d_i^*, \mathbf{d}_{-i}) \leq d_i^*(A_i(d_i^*, \mathbf{d}_{-i})) = 0$. On the other hand, $u_i(d_i, \mathbf{d}_{-i}) = d_i^*(A_i(\mathbf{d})) - P_i(\mathbf{d}) > 0$, so agent i can increase her utility by (mis)reporting declaration d_i (instead of d_i^*), violating incentive compatibility.

Finally, suppose there exist some $\mathbf{d}$ such that

$$A_i(\mathbf{d}) \notin \text{argmax}_S\{d_i(S) - p_i(S, \mathbf{d}_{-i})\}.$$

Then there is some S such that $d_i(S) - p_i(S, \mathbf{d}_{-i}) > d_i(A_i(\mathbf{d})) - p_i(A_i(\mathbf{d}), \mathbf{d}_{-i})$. In particular, $p_i(S, \mathbf{d}_{-i}) < \infty$, and hence, there exist some d_i' such that $A_i(d_i', \mathbf{d}_{-i}) \supseteq S$ and $p_i(A_i(d_i', \mathbf{d}_{-i}), \mathbf{d}_{-i}) = p_i(S, \mathbf{d}_{-i})$. We therefore conclude that

$$d_i(A_i(d_i', \mathbf{d}_{-i})) - P_i(d_i', \mathbf{d}_{-i}) > d_i(A_i(d_i, \mathbf{d}_{-i})) - P_i(d_i, \mathbf{d}_{-i}).$$

This implies that an agent with type d_i can increase her utility by (mis)reporting declaration d_i', violating incentive compatibility. □

2.3. Priority Algorithms

In this section, we review the priority algorithm framework [Borodin et al. 2003] 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, and so on. 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 all 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 type of behavior than is apparent in the definition of priority algorithms. Indeed, we view priority algorithms more generally as "greedy-like" or "myopic," in that the decision being made for the current input item being considered does not depend on input items not already processed. 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. For example, when the relevant decision is to either accept or reject a bid, a greedy priority algorithm must always accept every feasible bid (and, if there is flexibility in how to satisfy the bid, the allocation that maximizes social welfare must be chosen). With respect to our technical results, we note the difference between Theorems 3.1 and 4.1 which impose the greedy constraint and Theorems 3.5 and 3.6 which do not impose the greedy constraint. We note that many greedy CA algorithms in the literature are fixed order greedy priority algorithms.

.

3. TRUTHFUL PRIORITY ALGORITHMS

Lehmann et al. [2002] show that a greedy $O(\sqrt{m})$-approximation algorithm[8] 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 $v \leq v'$. 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,[9] but it can also apply to succinctly represented valuation functions.[10]

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.[11] Our main result is a lower bound on the approximation ratio,

[8]The Lehmann et al. [2002] algorithm will satisfy all models discussed in Section 3.

[9]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.

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

[11]Assume bidder i has already been allocated some set S_1. Then, when later 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 3.1, 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

achievable by a truthful greedy algorithm in the sets as items model. Theorem 3.1 implies a severe separation between the power of greedy algorithms and the power of truthful greedy algorithms. This separation holds even for the restricted case of s-CAs ($s \geq 2$), where bidders only desire sets having at most s objects. 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 3.1. *Suppose A is an incentive compatible and individually rational 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 the same 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 2.2, it must also maximize utility for agent i. The algorithm must therefore allocate the singleton which has the smaller minimal 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 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 minimal prices for singletons for one bidder is independent of the valuations of other bidders. Recall that $p_i(S, \mathbf{d}_{-i})$ is the minimal price for set S for bidder i, given $\mathbf{d}_{-i}$.

CLAIM 3.2. *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 $x = \alpha v^*$ for some sufficiently large α that we will set later (it will turn out that $\alpha \geq k^2$ will be sufficient). We will define various different possible valuation functions for bidder i: f, h, and g_c for all $c \in M$ and for sufficiently small $\epsilon > 0$.

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

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.

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

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 behavior 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 $x = \alpha v^*$ is possible by allocating $\{a, c\}$ to bidder i. Thus, as long as $\epsilon < v^*$ and $\alpha \geq k^2$, 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 2.2, 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 $I_4 \subseteq \mathcal{I}_1$. Since $(i, S, 0) \in I_4$, this item will be considered first by A on input I_4, 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 Claim 3.2. □

We can think of Claim 3.2 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 prices no less than the price of a for agent i. Note that $a \in T_i(a)$. 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.

Claim 3.3. *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 2.2, bidder i cannot be allocated set S, as required. □

Claim 3.3 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.

Claim 3.4. *There exists $a^* \in M$ such that $|R(a^*)| \geq k/2$.*

Proof. Note that for each i and any $x \in [m]$, the number of objects $a \in M$ for which $|T_i(a)| \geq x$ is at least $m - x + 1$; this follows because $\preceq_i$ defines an ordering over the objects in M, and all but the top $x - 1$ in this ordering must satisfy $|T_i(a)| \geq x$. This implies

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

Rearranging the order of summation, we also have

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

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

We are now ready to proceed with the proof of Theorem 3.1. Let $a^* \in M$ be the object from Claim 3.4. 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 Claim 3.3, 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 3.1. □

We believe that the greediness assumption of Theorem 3.1 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* $S_1, \ldots, S_\ell$ 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 [Nisan 2000]. 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. However, a feasible bid by agent i can be rejected and then a later bid by agent i can be accepted.

We now show that the greediness assumption from Theorem 3.1 can be removed when we consider priority algorithms in the elementary bids as items model.

THEOREM 3.5. *Suppose A is an incentive compatible and individually rational 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 claim that, for any i, it must be that $p_i(S, \mathbf{d}_{-i}) = 0$ for all $S \subseteq M$, whenever d_j is the 0 declaration for all $j \neq i$ (that is, only bidder i places non-zero bids). Indeed, suppose there exists S such that $p_i(S, \mathbf{d}_{-i}) = v > 0$. Let d_i be the single-minded declaration for set S with value $v/2$. Then, for each $T \supseteq S$, $d_i(T) = v/2 < p_i(S, \mathbf{d}_{-i})$. Since $p_i(S, \mathbf{d}_{-i})$ is the infimum of winning bids for sets containing S, we conclude that $A_i(d_i, \mathbf{d}_{-i}) \not\supseteq S$, and hence, $d_i(A_i(d_i, \mathbf{d}_{-i})) = 0$. Thus, A obtains a social welfare of 0 on input $(d_i, \mathbf{d}_{-i})$ when $v/2 > 0$ was possible, contradicting the supposed approximation ratio of A.

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, yielding an unbounded approximation factor: A achieves social welfare 0, whereas the optimal welfare is $1+\delta$.

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

We conclude that A must allocate T to bidder j. Let $I_4 \subseteq I_1$ be the input instance consisting of items $(j, T, 1)$ and $(j, M, 1+\delta)$. Then A will allocate T to bidder j in instance I_4. Recalling our earlier claim that $p_i(S, \mathbf{d}_{-i}) = 0$ for all $S \subseteq M$ whenever bidder i is the only non-zero bidder, we note that $d_j(T) - p_j(T, \mathbf{d}_{-j}) = 1 < 1+\delta = d_j(M) - p_j(M, \mathbf{d}_{-j})$. This, therefore, contradicts Theorem 2.2, which requires that the allocation to bidder j maximize the value of $d_i(S) - p_i(S, \mathbf{d}_{-i})$ over all sets S.

We therefore arrive at a contradiction in all cases, as required. $\square$

3.2. Bidders as Items

Roughly speaking, the lower bounds in Theorems 3.1 and 3.5 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. [2002] falls within this model, as does its (non-truthful) generalization to complex bidders [Lehmann et al. 2002], the primal-dual algorithm of Briest et al. [2011], and the (first) algorithm of Bartal et al. [2003] for multi-unit CAs. We now establish an inapproximation bound for truthful priority allocations that use bidders as items.

THEOREM 3.6. *Suppose A is an incentive compatible and individually rational priority algorithm for the (2-minded) CA problem that uses bidders as items. Then 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 $1 \le i \le n$ 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 and the iterative nature of priority algorithms. We define two sets of valuation functions, $\{g_1, \dots, g_k\}$ and $\{f_1, \dots, f_k\}$, that will be used in these input instances. The g_i will be defined so as to bound the minimal prices for the entire set of objects M. The f_i will be defined so as to force the algorithm into either forgoing many small values or choosing to award the grand bundle M to some agent who comes early in the iterative order adaptively chosen by the priority algorithm. The functions $g_1, \dots, g_k$ are straightforward: for each $1 \le i \le k$, 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 $f_1, \dots, f_k$ is more involved. Fix $i \in N$ and define $V'_{-i} := \{g_1, \dots, g_k\}^{n-1}$. 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 minimal price property, there is a minimal price $p_i(M, \mathbf{d}_{-i})$ for set M given this $\mathbf{d}_{-i}$.

CLAIM 3.7. *$p_i(M, \mathbf{d}_{-i}) \le kn$.*

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 minimal 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 Claim 3.7. □

We are now ready to define the valuations $f_1, \dots, f_k$. 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 minimal 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.

CLAIM 3.8. *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 $x > p_i(M, \mathbf{d}_{-i})$ from the definition of x. But now, from the definition of f_i,

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

for all $S \neq M$. Therefore, by the minimal pricing property (Theorem 2.2), A must allocate M to bidder i, completing the proof of Claim 3.8. □

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.

CLAIM 3.9. *There exists a labeling 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 $i - 1$. Recall that $I_{i-1} = \{(1, g_1), \ldots, (i-1, g_{i-1})\}$. Define $\mathcal{I}_i \subseteq \mathcal{I}'$ as follows:

$$\mathcal{I}_i := \{(j, v_j) \mid (j, v_j) \in \mathcal{I}', j > i - 1\}.$$

That is, $\mathcal{I}_i$ contains all of the items of $\mathcal{I}'$ that correspond to bidders not present in I_{i-1}. Note that if $I \subseteq \mathcal{I}'$ is a valid input instance such that $I_{i-1} \subseteq I$, then we must have $I \subseteq I_{i-1} \cup \mathcal{I}_i$.

Consider the execution of A on any valid input instance $I \subseteq I_{i-1} \cup \mathcal{I}_i$. By our induction hypothesis, the algorithm will first consider the items of I_{i-1} and allocate $\emptyset$ to each bidder $1, \ldots, i-1$. 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 $I \subseteq I_{i-1} \cup \mathcal{I}_i$ 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 Claim 3.8, so A must allocate M to bidder i. This is a contradiction. We conclude that this first case cannot occur.

Case 2: The first 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 $I \in I_{i-1} \cup \mathcal{I}_i$ 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 claim, completing the proof of Claim 3.9. □

Now suppose $I_{k/2-1}$ is the set from Claim 3.9 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 Claim 3.9, 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 3.6. □

4. TRUTHFUL SUBMODULAR PRIORITY AUCTIONS

Lehmann et al. [2006] 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. (We assume any fixed deterministic method to resolve ties among agents.) 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 4.1. *Any greedy priority algorithm for the combinatorial auction problem that uses objects as items is not incentive compatible. This holds even when there are only two bidders and when the bidders have submodular valuations.*

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 2.1, this valuation defines a price $\pi_2(S)$ for each subset $S \subseteq M$, and allocates to bidder 2 the utility-maximizing subset under these prices. Consider the prices for all subsets of size 2 and suppose, without loss of generality, that $\{a_1, a_2\}$ has the smallest price. That is, $\pi_2(\{a_1, a_2\}) \leq \pi_2(\{a_2, a_3\})$ and $\pi_2(\{a_1, a_2\}) \leq \pi_2(\{a_1, a_3\})$.

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 six 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 2.1,

$$v_2(\{a_1, a_3\}) - \pi_2(\{a_1, a_3\}) \geq v_2(\{a_1, a_2\}) - \pi_2(\{a_1, a_2\}).$$

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

$$\pi_2(\{a_1, a_2\}) > \pi_2(\{a_1, a_3\})$$

which contradicts the minimality of $\pi_2(\{a_1, a_2\})$.

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 any positive marginal value, without affecting the above construction. Such items can be considered at any time and can be allocated to any agent without changing the analysis. □

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. It would also be of some interest to close the gap between our $\frac{(1-\delta)k}{2}$ priority inapproximations and the naïve $k = \min\{n, m\}$ truthful approximation. 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 truthful sublinear 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 Angelopoulos and Borodin [2010]) 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 [Horn 2004] 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 [Borodin et al. 2011] (as a formalization of local ratio/primal dual algorithms with reverse delete [Bar-Noy et al. 2001]), 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 [Krysta 2005]. 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 [Bartal et al. 2003; Briest et al. 2011; Mu'alem and Nisan 2008] 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 [2011].

The results in this article 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.

ACKNOWLEDGMENTS

We thank S. Dobzinski, R. Gonen, and S. Micali for helpful discussions. We also thank the anonymous referees for many constructive suggestions.

REFERENCES

S. Angelopoulos and A. Borodin. 2010. Randomized priority algorithms. *Theoretical Comput. Sci.* 411 (2010), 2542–2558.

A. Archer, C. H. Papadimitriou, K. Talwar, and É. Tardos. 2003. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Math.* 1, 2 (2003), 129–150.

L. M. Ausubel and P. R. Milgrom. 2002. Ascending auctions with package bidding. *Frontiers of Theoretical Economics* 1, 1 (2002).

Y. Azar, I. Gamzu, and S. Gutner. 2010. Truthful unsplittable flow for large capacity networks. *ACM Trans. Algorithms* 6, 2 (2010).

M. Babaioff and L. Blumrosen. 2008. Computationally-feasible truthful auctions for convex bundles. *Games and Economic Behavior* 63, 2 (2008), 588–620.

A. Bar-Noy, R. Bar-Yehuda, A. Freund, J. Naor, and B. Schieber. 2001. A unified approach to approximating resource allocation and scheduling. *J. ACM* 48, 5 (2001), 1069–1090.

Y. Bartal, R. Gonen, and N. Nisan. 2003. Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'03)*. 72–87.

X. Bei and Z. Huang. 2011. Bayesian incentive compatibility via fractional assignments. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*. 720–733.

K. Bhawalkar and T. Roughgarden. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*. 700–709.

A. Borodin, D. Cashman, and A. Magen. 2011. How well can primal-dual and local-ratio algorithms perform? *ACM Trans. Algorithms* 7, 3 (2011), 29.

A. Borodin and B. Lucier. 2010. Greedy mechanism design for truthful combinatorial auctions. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP'10), Part I*. 90–101.

A. Borodin, M. N. Nielsen, and C. Rackoff. 2003. (Incremental) priority algorithms. *Algorithmica* 37, 4 (2003), 295–326.

P. Briest, P. Krysta, and B. Vöcking. 2011. Approximation techniques for utilitarian mechanism design. *SIAM J. Comput.* 40, 6 (2011), 1587–1622.

D. Buchfuhrer, S. Dughmi, H. Fu, R. Kleinberg, E. Mossel, C. Papadimitriou, M. Schapira, Y. Singer, and C. Umans. 2010. Inapproximability for VCG-based combinatorial auctions. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC'10)*. 518–536.

I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, and É. Tardos. 2015. Bounding the inefficiency of outcomes in generalized second price auctions. *J. Economic Theory* 156 (2015), 343–388. DOI:http://dx.doi.org/10.1016/j.jet.2014.04.010 Computer Science and Economic Theory.

B. Chandra and M. M. Halldórsson. 2001. Greedy local improvement and weighted set packing approximation. *J. Algorithms* 39, 2 (2001), 223–240.

S. Chawla, J. Hartline, D. Malec, and B. Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*. 311–320.

C. Chekuri and I. Gamzu. 2009. Truthful mechanisms via greedy iterative packing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, (APPROX'09), and 13th International Workshop (RANDOM'09)*. 56–69.

G. Christodoulou, A. Kovács, and Michael Schapira. 2008. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP'08) Part I: Tack A: Algorithms, Automata, Complexity, and Games*. 820–832.

P. Cramton. 2002. Spectrum auctions. In *Handbook of Telecommunications Economics*, M. Cave, S. Majumdar, and I. Vogelsang (Eds.). Elsevier Science B. V., Chapter 14, 605–639.

P. Cramton, Y. Shoham, and R. Steinberg (Eds.). 2005. *Combinatorial Auctions*. MIT Press, Cambridge, MA.

A. Daniely, M. Schapira, and G. Shahaf. 2015. Inapproximability of truthful mechanisms via generalizations of the VC dimension. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC'15)*. 401–408.

P. Dhangwatnotai, S. Dobzinski, S. Dughmi, and T. Roughgarden. 2011. Truthful approximation schemes for single-parameter agents. *SIAM J. Comput.* 40, 3 (2011), 915–933.

S. Dobzinski. 2007. 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'07)*. 89–103.

S. Dobzinski. 2011. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*.

S. Dobzinski and S. Dughmi. 2013. On the power of randomization in algorithmic mechanism design. *SIAM J. Comput.* 42, 6 (2013), 2287–2304.

S. Dobzinski and N. Nisan. 2011. Limitations of VCG-based mechanisms. *Combinatorica* 31, 4 (2011), 379–396.

S. Dobzinski, N. Nisan, and M. Schapira. 2010. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.* 35, 1 (2010), 1–13.

S. Dobzinski, N. Nisan, and M. Schapira. 2012. Truthful randomized mechanisms for combinatorial auctions. *J. Comput. Syst. Sci.* 78, 1 (2012), 15–25.

S. Dobzinski and M. Schapira. 2006. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'06).* 1064–1073.

S. Dobzinski and M. Sundararajan. 2008. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *Proceedings 9th ACM Conference on Electronic Commerce (EC'08).* 38–47.

S. Dobzinski and J. Vondrák. 2012. The computational complexity of truthfulness in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC'12).* 405–422.

P. Dütting, V. Gkatzelis, and T. Roughgarden. 2014. The performance of deferred-acceptance auctions. In *Proceedings of the ACM Conference on Economics and Computation (EC'14).* 187–204.

B. Edelman, M. Ostrovsky, and M. Schwarz. 2005. 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*.

M. Feldman, H. Fu, N. Gravin, and B. Lucier. 2013. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. ACM, New York, NY, 201–210.

R. Gonen and D. Lehmann. 2000. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of the ACM Conference on Economics and Computation (EC'10).* 13–20.

V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. 2005. On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05).* 1164–1173.

J. Hartline, R. Kleinberg, and A. Malekian. 2011. Bayesian incentive compatibility via matchings. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'11).* 734–747.

J. Hartline and B. Lucier. 2010. Bayesian algorithmic mechanism design. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10).* 301–310.

J. Hastad. 1999. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica* 182 (1999), 105–142.

R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. 2004. Bundling equilibrium in combinatiorial auctions. *Games and Economic Behavior* 47 (2004), 104–123.

S. Horn. 2004. *One-Pass algorithms with revocable acceptances for job interval selection*. MSC thesis. University of Toronto.

T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. 2013. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Proceeddings of the 21st Annual European Symposium on Algorithms, Sophia Antipolis, France, September 2–4, 2013*. 589–600.

S. Khot, R. J. Lipton, E. Markakis, and A. Mehta. 2008. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica* 52, 1 (2008), 3–18.

P. Krysta. 2005. 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*. 615–627.

R. Lavi, A. Mu'alem, and N. Nisan. 2003. Towards a characterization of truthful combinatorial auctions. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, Proceedings*. 574–583.

R. Lavi and N. Nisan. 2015. Online ascending auctions for gradually expiring items. *J. Economic Theory* 156 (2015), 45–76.

R. Lavi and C. Swamy. 2011. Truthful and near-optimal mechanism design via linear programming. *J. ACM* 58, 6 (2011), 25.

B. Lehmann, D. J. Lehmann, and N. Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2 (2006), 270–296.

D. J. Lehmann, L. O'Callaghan, and Y. Shoham. 2002. Truth revelation in approximately efficient combinatorial auctions. *J. ACM* 49, 5 (2002), 577–602.

B. Lucier. 2010. Beyond equilibria: Mechanisms for repeated combinatorial auctions. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*. 166–177.

B. Lucier and A. Borodin. 2010. Price of anarchy for greedy auctions. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, January 17-19, 2010*. 537–553.

M. Mahdian and A. Saberi. 2006. Multi-unit auctions with unknown supply. In *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006), Ann Arbor, Michigan, June 11-15, 2006*. 243–249.

A. Mu'alem and N. Nisan. 2008. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior.* 64 (2008), 612–631. Issue 2.

N. Nisan. 2000. Bidding and allocation in combinatorial auctions. In *EC*. 1–12.

N. Nisan and A. Ronen. 2001. Algorithmic mechanism design. *Games and Economic Behavior* 35, 1–2 (2001), 166–196.

C. Papadimitriou, M. Schapira, and Y. Singer. 2008. On the hardness of being truthful. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA*. 250–259.

S. J. Rassenti, V. L. Smith, and R. L. Bulfin. 1982. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics.* 13 (1982), 402–417.

V. Syrgkanis and É. Tardos. 2013. Composable and efficient mechanisms. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, June 1–4, 2013*. 211–220.

D. Zuckerman. 2007. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing* 3, 1 (2007), 103–128.

Received January 2015; revised August 2015; accepted November 2015