

Price of Anarchy for Greedy Auctions

B. Lucier*

A. Borodin†

July 6, 2009

Abstract

We consider mechanisms for utilitarian combinatorial allocation problems, where agents are not assumed to be single-minded. This class of problems includes combinatorial auctions, multi-unit auctions, unsplittable flow problems, profit-maximizing scheduling, and others. We study the price of anarchy for such mechanisms, which is a bound on the approximation ratio obtained at any mixed Nash equilibrium. We demonstrate that a broad class of greedy approximation algorithms can be implemented as mechanisms for which the price of anarchy nearly matches the performance of the original algorithm. This is true even in Bayesian settings, where the agents' valuations are drawn from publicly known distributions. Furthermore, for a rich subclass of allocation problems, pure Nash equilibria are guaranteed to exist for these mechanisms. For many problems, the approximation factors obtained at equilibrium improve upon the best known results for deterministic truthful mechanisms. In particular, we exhibit a simple deterministic mechanism for the general combinatorial auction with $O(\sqrt{m})$ price of anarchy.

*Dept of Computer Science, University of Toronto, blucier@cs.toronto.edu.

†Dept of Computer Science, University of Toronto, bor@cs.toronto.edu.

1 Introduction

The field of algorithmic mechanism design lies at the intersection of game-theoretic and computational concerns for interactive systems. The marriage of these two settings has spawned a fruitful line of research aimed at answering a primary question: can any computationally efficient algorithm be converted into a computationally efficient mechanism for selfish agents? For utilitarian social choice functions, the celebrated Vickrey-Clarke-Groves (VCG) mechanism addresses game-theoretic issues in a strong sense: in the absence of collusion, it induces full cooperation (ie. truth-telling) from all agents as a dominant strategy. However, the VCG mechanism requires that the underlying welfare-optimization problem be solved exactly, and is therefore ill-suited to computationally intractable problems. The standard computational answer to such problems is the development of approximation algorithms, but the VCG mechanism does not (in general) retain its truthfulness when applied to such approximate solutions [23].

The incompatibility between approximations and standard mechanism design techniques has motivated the search for new algorithms that are specially tailored for selfish agents. This search has focused primarily on the special case of incentive compatible (IC) algorithms, which are precisely those that can be implemented as truthful mechanisms. While this venture has been largely successful in settings where agent preferences are single-dimensional [21, 7, 19, 1], general settings have proven more difficult. Indeed, it has been shown that the approximation ratios achievable by IC polytime deterministic algorithms and their non-IC counterparts exhibit a large asymptotic gap for some problems [25].

We consider an alternative to truthfulness in dominant strategies: (mixed) Bayesian Nash equilibrium (BNE). We suppose that the types of the agents are drawn from known (not necessarily identical) distributions. After types are chosen, each agent applies a strategy (or a distribution over strategies) that maximizes his utility, given the distribution of the strategies of the other agents. In a BNE, no agent will have incentive to unilaterally deviate from this equilibrium. The BNE solution concept is of primary importance in the economics and implementation theory literature; see Jackson [16] for a survey. We pose the basic question: can a given approximation algorithm be converted into a mechanism that preserves its approximation ratio at every BNE? We show that for a broad class of non-IC greedy algorithms, the answer is *yes*. Note that the BNE solution concept is not, strictly speaking, a relaxation of dominant strategy truthfulness. There exist truthful mechanisms whose approximation ratios are not preserved at all Nash equilibria, such as the famous Vickrey auction [8].

We restrict our attention to combinatorial allocation problems, where the goal is to assign m objects to n agents in such a way that the overall social welfare is maximized. We allow arbitrary feasibility constraints to determine which allocations are permitted, and we do not restrict agents to be single-minded. This class includes the well-studied combinatorial auction (CA) problem, as well as multi-unit CAs, the unsplittable flow problem, and many others. We then consider a broad class of “greedy algorithms” (explicitly described below) for approximately solving such allocation problems. Such algorithms are not generally incentive compatible [19]. Our first result is that if such a greedy algorithm is paired with a first-price payment scheme (i.e. each agent pays his declared value for the set he receives), the resulting auction nearly preserves the original algorithm’s approximation ratio at every BNE.

Theorem: Any greedy c -approximation algorithm for a combinatorial allocation problem can be implemented as a deterministic mechanism that achieves a $(c+O(\log c))$ approximation at every Bayesian Nash equilibrium. There exist examples in which the first-price mechanism can have a $(c + \Omega(\log c))$ approximation at a mixed Nash equilibrium.

Are there implementations of a c -approximate greedy algorithm for which the price of anarchy is not $c + \theta(\log c)$, but rather c ? As a step towards resolving this question, we consider an alternative mechanism that charges so-called *critical prices*. In a critical-price payment scheme, a winning agent pays the smallest amount he could have bid on the set he receives and still won it. As has been noted elsewhere [8], mechanisms of this form can suffer from unnatural problems at equilibrium: an agent may have incentive to greatly over-represent his values, hoping that no other agent makes large bids. Indeed, we construct examples in which an agent might have a strict preference for doing so. This possibility of over-bidding can result in equilibria with poor social welfare. However, these bidding strategies are inherently risky: depending on the bids of other agents, an over-bidding agent may end up with negative utility. We eliminate this risky behaviour by making two concessions. First, we adopt

a common assumption from mechanism design, that the agents are *ex-post individually-rational*, which means loosely that no agent bids in such a way that he might obtain negative utility. Second, we perturb the payment rule of the mechanism so that, with positive probability, an agent may have to pay his winning bid. These two factors imply that no agent will declare more than his true value for a set.

Theorem: If agents are ex-post individually rational, then any greedy c -approximation algorithm for a combinatorial allocation problem can be implemented as an auction that achieves a $(c+1)$ approximation at every Bayesian Nash equilibrium.

For certain algorithms, we can tighten our result even further so that the price of anarchy is c . This includes algorithms that are $(c-1)$ -approximate when agents are single-minded, as well as algorithms that are symmetric with respect to the agents and objects.

Our perturbed critical-price mechanism is not deterministic. However, we note a difference in the use of randomness here versus other randomized truthful mechanisms in the literature [10, 17]. We are not implementing a randomized allocation rule; rather, our allocation rule is deterministic, but our payment scheme includes an alternative that is applied with arbitrarily low probability.

We next turn our attention to pure Nash equilibria in the full-information (ie. non-Bayesian) setting. We show that the price of anarchy for the deterministic first-price mechanism is improved when restricted to pure equilibria. However, the existence of an equilibrium in pure strategies is not guaranteed in general. To avoid well-known issues in which equilibria do not exist because of infinitesimal improvements, we consider the solution concept of ϵ -Nash equilibrium, whereby agents are indifferent about outcomes whose utilities differ by at most ϵ . Even with this discretization, we demonstrate that for some allocation problems and greedy algorithms, neither the first-price mechanism nor the critical-price mechanism will necessarily have a pure ϵ -Nash equilibrium.

We consider two important special cases for which we can guarantee the existence of pure equilibria (and retain the approximation ratio). First, we present the *blocking allocation problems* (in which any agent can be allocated any given pair of objects), for which a pure equilibrium is guaranteed to exist for a broad class of our greedy allocation rules (those that are *non-adaptive* and *continuous*, to be formally defined later) paired with a first-price payment scheme.

Second, we can remove the restriction on the problem specification, and instead focus on the *standard greedy algorithm* that makes assignments in order of their value. This particular algorithm is well-studied; it is known to be k -approximate (respectively, $(k+1)$ -approximate) for linear (resp. submodular) functions on a k -independence system (e.g. for matroids, $k=1$). We show that the (deterministic) critical-price mechanism for a standard greedy algorithm always has a pure equilibrium. This mechanism retains the approximation ratio of the original algorithm at all equilibria, *assuming that agents do not over-bid*. We motivate this assumption (which is stronger than ex-post individual rationality) by noting that for each agent, when the declarations of the other agents are fixed (as in a pure equilibrium), not over-bidding forms a weakly dominant set of strategies.

We conclude by showing how a combination of a greedy allocation rule with a rule that allocates all objects to a single bidder can be implemented while retaining the price of anarchy and equilibrium existence properties discussed in this paper.

1.1 Related Work

The most prominently studied allocation problem that falls into our framework is the combinatorial auction problem. Hastad's [13] result shows that it is NP-hard to approximate CAs to within $\Omega(\sqrt{m})$ even for succinctly representable valuation functions. An alternative bound by Nisan [22] shows that exponential communication is necessary to approximate general CAs to within $\Omega(\sqrt{m})$, independent of complexity assumptions. The best known deterministic truthful mechanism for CAs with general valuations attains an approximation ratio of $O(\frac{m}{\sqrt{\log m}})$ [14]. A randomized $O(\sqrt{m})$ -approximate mechanism that is truthful in expectation was given by Lavi and Swamy [17]. Dobzinski, Nisan and Schapira [10] then gave a universally truthful randomized mechanism that attains an $O(\sqrt{m})$ approximation.

Many variations and restrictions on combinatorial auctions have been considered in the literature. Bartal et. al. [4] give a truthful $O(m^{\frac{1}{B-2}})$ mechanism for multi-unit combinatorial auctions with B copies of each object, for all $B \geq 3$. Dobzinski and Nisan [9] construct a truthful 2-approximate mechanism

for multi-unit auctions (ie. having many copies of just a single object), and a truthful PTAS when additionally each declaration can be represented as the maximum of k single-minded desires. Many other problems have truthful mechanisms ([19, 21, 7]) when bidders are restricted to being single-minded.

The BNE solution concept was recently applied to submodular combinatorial auctions [8], where it was shown that a randomized mechanism can attain a 2-approximation at any mixed equilibrium assuming that bidders are ex-post individually rational. Pure equilibria of first-price mechanisms have also been studied for path procurement auctions [15]. Performance at Nash equilibrium has been extensively studied in the economics literature (see Jackson [16] for a survey) and recently in work on Internet advertising slot auctions [26, 11]. In that line of research the goal is often revenue maximization, rather than social welfare maximization, and traditionally one wishes to implement a particular optimal allocation rule, rather than guarantee a given approximation ratio.

The price of anarchy has been well-studied in many different game contexts, though generally for settings where agents choose their outcomes directly rather than through a mechanism. The literature includes many refinements of these concepts, such as convergence of potential games and price of total anarchy. See chapters 17-21 of [24], and references therein.

The properties of greedy algorithms have been extensively studied. Borodin et al [6] introduced the notion of priority algorithms as a model for greedy algorithms, and studied their power in solving various approximation problems. Monotone greedy algorithms for combinatorial auctions were studied by Mu'alem and Nisan [21] and subsequently by Briest, Krysta, and Vocking [7], resulting in the development of new incentive compatible algorithms for single-minded bidders. Gonen and Lehmann [12] gave lower bounds on the power of greedy mechanisms to solve combinatorial auctions with general bidders.

2 Model and Definitions

2.1 Feasible Allocation Problems

We consider a setting in which there are n agents and a set M of m objects. An *allocation* to agent i is a subset $X_i \subseteq M$. A *valuation function* $v : 2^M \rightarrow \mathbf{R}$ assigns a value to each allocation. We assume that valuation functions are monotone, meaning $v(S) \leq v(T)$ for all $S \subseteq T \subseteq M$, and normalized so that $v(\emptyset) = 0$. 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 profile* is a vector of n valuation functions, one for each agent. In general we will use boldface to represent vectors, subscript i to denote the i th component, and subscript $-i$ to denote all components except i , so that $\mathbf{v} = (v_i, \mathbf{v}_{-i})$. An *allocation profile* is a vector of n allocations. A *combinatorial allocation problem* is defined by a set of *feasible allocations*, which is the set of permitted allocation profiles. An *allocation rule* \mathbf{X} assigns to each valuation profile \mathbf{v} a feasible outcome $\mathbf{X}(\mathbf{v})$; we write $X_i(\mathbf{v})$ for the allocation to agent i . An algorithm \mathcal{A} for a given combinatorial allocation problem is an implementation of an allocation rule. We write $\mathbf{X}^{\mathcal{A}}$ for the rule implemented by \mathcal{A} ; we will drop the superscript \mathcal{A} when it is clear from context.

Each agent $i \in [n]$ has a private valuation function t_i , his *type*, which defines the value he attributes to each allocation¹. The *social welfare* obtained by allocation profile \mathbf{X} , given type profile \mathbf{t} , is $SW(\mathbf{X}, \mathbf{t}) = \sum_i t_i(X_i)$. We write $SW_{opt}(\mathbf{t})$ for $\max_{\mathbf{X}} \{SW(\mathbf{X}, \mathbf{t})\}$ and say that algorithm \mathcal{A} is a c approximation algorithm if $SW(\mathbf{X}^{\mathcal{A}}(\mathbf{t}), \mathbf{t}) \geq \frac{1}{c} SW_{opt}(\mathbf{t})$ for all \mathbf{t} .

A payment rule P assigns a vector of n payments to each valuation profile. A *direct revelation mechanism* \mathcal{M} is composed of an allocation rule \mathbf{X} and a payment rule P . The mechanism proceeds by eliciting a valuation profile \mathbf{d} from each of the agents, called the *declarations*. It then applies the allocation and payment rules to \mathbf{d} to obtain an allocation and payment for each agent. Crucially, the agents may not declare their true types; that is, it may be that $\mathbf{t} \neq \mathbf{d}$.

Fix some mechanism \mathcal{M} . The utility of agent i , given declaration profile \mathbf{d} , is $u_i(\mathbf{d}) = t_i(X_i(\mathbf{d})) - P_i(\mathbf{d})$. Declaration profile \mathbf{d} forms a *pure Nash equilibrium* if, for all $i \in [n]$ and all d_i' , $u_i(d_i, \mathbf{d}_{-i}) \geq u_i(d_i', \mathbf{d}_{-i})$. That is, no one player can obtain a higher utility by deviating from declaration \mathbf{d} . An ϵ -Nash equilibrium for some $\epsilon > 0$ is a profile \mathbf{d} such that $u_i(d_i, \mathbf{d}_{-i}) \geq u_i(d_i', \mathbf{d}_{-i}) - \epsilon$.

¹We present our results in a non-Bayesian setting. See Appendix A for an extension to Bayesian types.

Given a sequence of probability distributions $\omega_1, \dots, \omega_n$ over declarations, and any function f over the space of declaration profiles, we will write $E_{\mathbf{d} \sim \omega}[f(\mathbf{d})]$ for the expected value of f over declarations chosen according to the product distribution $\omega = \omega_1 \times \dots \times \omega_n$. Product distribution ω is a *mixed Nash equilibrium* if, for each $i \in [n]$ and all distributions ω'_i ,

$$E_{(d_i, \mathbf{d}_{-i}) \sim (\omega_i, \omega_{-i})}[u_i(d_i, \mathbf{d}_{-i})] \geq E_{(d'_i, \mathbf{d}_{-i}) \sim (\omega'_i, \omega_{-i})}[u_i(d'_i, \mathbf{d}_{-i})]. \quad (1)$$

That is, the distribution maximizes the expected utility for each agent, given the distributions of the others. The *price of anarchy of \mathcal{M} in mixed and pure strategies* are defined as

$$PoA_{mixed} = \sup_{\mathbf{t}, \omega} \frac{SW_{opt}(\mathbf{t})}{E_{\mathbf{d} \sim \omega}[SW(X(\mathbf{d}), \mathbf{t})]} \quad PoA_{pure} = \sup_{\mathbf{t}, \mathbf{d}} \frac{SW_{opt}(\mathbf{t})}{SW(X(\mathbf{d}), \mathbf{t})}$$

where the supremums are over all type profiles \mathbf{t} and all mixed Nash equilibria ω (respectively, all pure Nash equilibria \mathbf{d}) for \mathbf{t} . Whenever a pure Nash exists, we have $PoA_{pure} \leq PoA_{mixed}$.

2.2 Greedy Allocation Rules

We describe a special type of allocation rule, which we will refer to as a *greedy allocation rule*. These are motivated by the monotone greedy algorithms of Mu'alem and Nisan [21], extended to be adaptive. We begin with some definitions. A *partial allocation profile* is a sequence of allocations, one for each i in some subset S of $[n]$. A partial allocation profile is *feasible* if there is some feasible allocation profile that extends it. Given a partial allocation profile for subset S , some $i \notin S$, and allocation X_i , we say X_i is a *feasible allocation for i given S* if the partial allocation remains feasible when X_i is added to it.

A *monotone priority function* is a function $r : [n] \times 2^M \times \mathbf{R} \rightarrow \mathbf{R}$. $r(i, S, v)$ is the priority of allocating $S \subseteq M$ to player i when $v_i(S) = v$; we require for r to be monotone non-decreasing in v and monotone non-increasing in S with respect to inclusion. We consider two types of greedy allocation rules. A *non-adaptive greedy allocation rule* is an allocation algorithm of the following form:

1. Fix a monotone priority function r . Let $N = [n]$.
2. Repeat until $N = \emptyset$:
3. Choose $i \in N$ and $S \subseteq M$ that maximizes $r(i, S, d_i(S))$ over all feasible allocations
4. Set $X_i = S$; remove player i from N
5. return X_1, \dots, X_n

A non-adaptive algorithm fixes a single priority function that is used throughout its execution. By contrast, an *adaptive greedy allocation rule* can change its priority function on each iteration, depending on the partial allocation formed on the previous iterations. Note that our definition of greedy allocation rules explicitly allows only a single allocation to each agent. This is in contrast to a very different type of “greedy-like” allocation rule, in which one iterates over the objects and the allocation to each agent is built up incrementally (e.g. for submodular combinatorial auctions [18]). Such incremental allocation rules are not covered by our results; we leave open the BNE analysis of their implementations.

We now mention a few examples of combinatorial allocation problems and greedy allocation rules. (See also section J.) The general combinatorial auction problem is defined by the feasibility constraint that no two allocations can intersect. Lehmann et al [19] show that the greedy allocation rule with $r(i, S, v) = \frac{v}{\sqrt{|S|}}$ achieves an $O(\sqrt{m})$ approximation ratio for CAs. The k -CA problem has the feasibility constraint that no two allocations can intersect, and additionally no allocated set can have size greater than k . The non-adaptive *standard greedy allocation rule* defined by $r(i, S, v) = v$ attains a $(k + 1)$ approximation. In the multi-unit CA problem, we think of there being B copies of each object for some $B \geq 1$. This problem is defined in our framework by the feasibility constraint that no more than B allocated sets contain any given object. A greedy algorithm attains an $O(m^{\frac{1}{B+1}})$ approximation when bidders are assumed to be single-minded [7]².

²We note that the associated algorithm for general bidders, GREEDY-2 in the same paper [7], is not a greedy algorithm as we define it, due to a correction step at its end. The same holds for the truthful $O(m^{\frac{1}{B-2}})$ approximation algorithm due to Bartal et al [4].

In a unit-job profit-maximizing scheduling problem, the objects are unit length intervals within some prescribed release time and deadline. The standard greedy allocation rule yields a 3-approximation for this problem [20]. In the unsplittable flow problem (UFP), we are given an undirected graph with edge capacities. The objects are the edges, and each valuation function is such that agent i has some value $v(s, t)$ for being given a path from s to t . An allocation is feasible if no edge is included in more allocated paths than its edge capacity allows. Let B be the minimum edge capacity. A primal-dual algorithm, which is an adaptive greedy allocation rule, obtains an $O(m^{1/(B-1)})$ approximation for any $B > 1$ [7].

We note that many of the algorithms above are known to be incentive compatible given that agents are single-minded. However, it is known that for general valuations and some problems (eg. CAs), no incentive compatible greedy algorithms (of the form we consider) can obtain non-trivial approximation ratios [5]. We do not assume that agents are single-minded, and consider the game-theoretic properties of these (non-incentive-compatible) algorithms when each agent has a multitude of preferences.

2.3 Payment Rules

Given an allocation rule \mathbf{X} , agent i , declaration profile \mathbf{d}_{-i} , and set S , the *critical price* $\theta_i(S, \mathbf{d}_{-i})$ for set S is the minimum amount that agent i could bid on set S and win S . That is,

$$\theta_i(S, \mathbf{d}_{-i}) = \inf\{v : \exists d_i, d_i(S) = v, X_i(d_i', \mathbf{d}_{-i}) = S\}.$$

The critical payment scheme sets $P_i(\mathbf{d}) = \theta_i(X_i(\mathbf{d}), \mathbf{d}_{-i})$, so each agent pays the critical price for the set she receives. We discuss implementation issues for this payment scheme in Section G. The first-price payment scheme sets $P_i(\mathbf{d}) = d_i(X_i(\mathbf{d}))$, so each agent pays her declared value for the set she receives.

Given an allocation rule \mathcal{A} , we will write $\mathcal{M}_1(\mathcal{A})$ to denote the direct revelation mechanism that applies allocation rule \mathcal{A} and the first-price payment scheme. We call this the *first-price mechanism for \mathcal{A}* . The *critical-price mechanism for \mathcal{A}* , $\mathcal{M}_{crit}(\mathcal{A})$, instead applies the critical price payment scheme. Given $\mu \in [0, 1]$, we denote by $\mathcal{M}_\mu(\mathcal{A})$ the mechanism that applies allocation rule \mathcal{A} and the following randomized payment scheme: a biased coin that lands heads with probability μ is flipped. If the coin lands heads, the first-price payment scheme is used, otherwise the critical-price payment scheme is used.

3 Bayesian Mixed Nash Equilibria

In this section we demonstrate how to implement a greedy allocation rule so that, at any Bayesian Nash Equilibrium of the resulting mechanism, the approximation ratio of the greedy rule is nearly preserved. We first analyze the price of anarchy for the first-price mechanism $\mathcal{M}_1(\mathcal{A})$. We then show how to improve our result using mechanism $\mathcal{M}_\mu(\mathcal{A})$, under the assumption that bidders are ex-post individually rational.

3.1 Properties of Greedy Algorithms

In all of the following, we assume that \mathcal{A} is an adaptive greedy algorithm for an arbitrary combinatorial allocation problem. We let \mathbf{X} denote the allocation rule for \mathcal{A} . We will write d_i^0 to mean the zero valuation as declared by agent i . The following lemma encapsulates the primary property of greedy algorithms that we will use: the critical prices and allocation for a given bidder depend only on the sets won by the other players, and their declared values for those sets.

Lemma 3.1. *Choose $i \in [n]$, and suppose \mathbf{d}_{-i} and \mathbf{d}'_{-i} are such that $\mathbf{X}(d_i^0, \mathbf{d}_{-i}) = \mathbf{X}(d_i^0, \mathbf{d}'_{-i})$ and $d_j(X_j(d_i^0, \mathbf{d}_{-i})) = d_j(X_j(d_i^0, \mathbf{d}'_{-i}))$ for all $j \neq i$. Then $\theta_i(S, \mathbf{d}_{-i}) = \theta_i(S, \mathbf{d}'_{-i})$ for all $S \subseteq M$, and $X_i(d_i, \mathbf{d}_{-i}) = X_i(d_i, \mathbf{d}'_{-i})$ for all d_i .*

Proof. Choose $i \in [n]$ and $S \subseteq M$. Consider the iterations of \mathcal{A} on input (d_i^0, \mathbf{d}_{-i}) ; there is some k such that allocating S to i is feasible for the first k iterations, and infeasible for all subsequent iterations (note that we may have $k = n$). For each iteration ℓ up to k , let v_ℓ be the minimal value such that tuple (i, S, v_ℓ) would appear first in the ranking for that iteration (or ∞ if no such value exists). Let $v = \min_\ell \{v_\ell\}$. If agent i makes a single-minded declaration for S at value v or more, he will be allocated set S ; and for any d_i with $d_i(S) < v$, $X_i(d_i, \mathbf{d}_{-i}) \neq S$. Thus $\theta_i(S, \mathbf{d}_{-i}) = v$. By the same argument, we

have $\theta_i(S, \mathbf{d}'_{-i}) = v$, since on inputs (d_i^0, \mathbf{d}_{-i}) and $(d_i^0, \mathbf{d}'_{-i})$ \mathcal{A} allocates the same sets (by assumption), and in the same order (since $d_j(X_j(d_i^0, \mathbf{d}_{-i})) = d_j(X_j(d_i^0, \mathbf{d}'_{-i}))$ for all $j \neq i$).

Of all sets S such that $d_i(S) \geq \theta_i(S, \mathbf{d}_{-i})$, agent i will be allocated the one which has highest ranking in the earliest iteration. As critical prices for agent i are identical on inputs (d_i, \mathbf{d}_{-i}) and (d_i, \mathbf{d}'_{-i}) , as is the behaviour of \mathcal{A} on each iteration before any set is allocated to agent i , the allocation to agent i will be the same on declarations (d_i, \mathbf{d}_{-i}) and (d_i, \mathbf{d}'_{-i}) , as required. \square

We now show that if \mathcal{A} always obtains a c -approximation to the optimal total welfare (with respect to the input valuation profile \mathbf{d}), then it is also a c -approximation to the sum of the critical prices of the optimal allocation profile.

Lemma 3.2. *If \mathcal{A} is a c -approximation, then for any declaration profile \mathbf{d} and allocation profile \mathbf{A} , $\sum_{i \in [n]} d_i(X_i^{\mathcal{A}}(\mathbf{d})) \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(A_i, \mathbf{d}_{-i})$.*

Proof. Choose any $\epsilon > 0$. Let d_i' be the single-minded declaration for set A_i at value $\theta_i(A_i, \mathbf{d}_{-i}) - \epsilon$. Let d_i^* be the pointwise maximum of d_i' and d_i . Repeated application of Lemma 3.1 for each agent implies that the allocations of \mathbf{d} and \mathbf{d}^* are identical. Since \mathcal{A} is a c -approximation, we conclude that $SW(\mathbf{X}^{\mathcal{A}}(\mathbf{d}), \mathbf{d}) = SW(\mathbf{X}^{\mathcal{A}}(\mathbf{d}^*), \mathbf{d}^*) \geq \frac{1}{c} SW(\mathbf{A}, \mathbf{d}^*) - n\epsilon \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(A_i, \mathbf{d}_{-i}) - n\epsilon$. The result follows by taking the limit as $\epsilon \rightarrow 0$. \square

3.2 Greedy First-Price Mechanisms

We now consider the price of anarchy of the first-price mechanism $\mathcal{M}_1(\mathcal{A})$ given a greedy allocation rule \mathcal{A} . We first note that, at equilibrium, no player ever overbids on a set that he may be allocated with positive probability.

Lemma 3.3. *Suppose $\omega = (\omega_1, \dots, \omega_n)$ is a mixed NE for mechanism $\mathcal{M}_1(\mathcal{A})$, given fixed types \mathbf{t} . If d_i has positive probability in ω_i , and $\Pr_{\mathbf{d}_{-i} \sim \omega_{-i}}[X_i(d_i, \mathbf{d}_{-i}) = S] > 0$, then $d_i(S) \leq t_i(S)$.*

We are now ready to prove an upper bound on the price of anarchy of $\mathcal{M}_1(\mathcal{A})$.

Theorem 3.4. *The expected welfare at any mixed Bayesian Nash equilibrium of $\mathcal{M}_1(\mathcal{A})$ is a $c + O(\log c)$ approximation to the optimal welfare.*

Proof. The proof of the stated bound is somewhat technical and we will instead prove the simpler result that any mixed Nash equilibrium is a $2(c+1)$ approximation. The improvement to $c + O(\log c)$ and the extension to Bayesian settings are deferred to Appendix A.

Fix a true type profile \mathbf{t} . Let $\omega = (\omega_1, \dots, \omega_n)$ be a probability distribution on the set of all possible declarations, and suppose that ω forms a mixed Nash equilibrium. Let $\mathbf{A} = A_1, \dots, A_n$ denote an optimal allocation for \mathbf{t} . Let \mathbf{d} be a declaration profile in the support of ω . By Lemma 3.3, $d_i(S) \leq t_i(S)$ for all i and S such that $S = X_i(\mathbf{d})$. Lemma 3.2 then implies that

$$\sum_{i \in [n]} t_i(X_i(\mathbf{d})) \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(A_i, \mathbf{d}_{-i}) = \frac{1}{2c} \sum_{i \in [n]} t_i(A_i) - \frac{1}{2c} \sum_{i \in [n]} [t_i(A_i) - 2\theta_i(A_i, \mathbf{d}_{-i})].$$

Summing over all declaration profiles, with respect to probability function ω , we have that

$$\begin{aligned} \sum_{\mathbf{d}} \omega(\mathbf{d}) \sum_i t_i(X_i(\mathbf{d})) &\geq \sum_{\mathbf{d}} \omega(\mathbf{d}) \frac{1}{2c} \sum_i t_i(A_i) \\ &\quad - \sum_{\mathbf{d}} \omega(\mathbf{d}) \frac{1}{2c} \sum_i [t_i(A_i) - 2\theta_i(A_i, \mathbf{d}_{-i})] \\ &= \frac{1}{2c} \sum_i t_i(A_i) - \frac{1}{2c} \sum_i \sum_{\mathbf{d}} \omega(\mathbf{d}) [t_i(A_i) - 2\theta_i(A_i, \mathbf{d}_{-i})]. \end{aligned} \tag{2}$$

For each $i \in [n]$, let $y_i = 2E_{\mathbf{d}_{-i} \sim \omega_{-i}}[\theta_i(A_i, \mathbf{d}_{-i})]$ be twice the expected critical price for set A_i . Let d_i' be the single-minded declaration for set A_i at value y_i . If agent i were to declare d_i' with probability 1, he would obtain set A_i and pay y_i whenever $\theta_i(A_i, \mathbf{d}_{-i}) < y_i$. We therefore have

$$\begin{aligned} E_{\mathbf{d}_{-i} \sim \omega_{-i}}[t_i(X_i(d_i'', \mathbf{d}_{-i}))] &= (Pr_{\mathbf{d}_{-i} \sim \omega_{-i}}[\theta_i(A_i, \mathbf{d}_{-i}) \leq y_i])(t_i(A_i) - y_i) \\ &\geq \frac{1}{2}(t_i(A_i) - 2E_{\mathbf{d}'_{-i} \sim \omega_{-i}}[\theta_i(A_i, \mathbf{d}'_{-i})]) \\ &= \frac{1}{2} \sum_{\mathbf{d}} \omega(\mathbf{d})[t_i(A_i) - 2\theta_i(A_i, \mathbf{d}_{-i})] \end{aligned}$$

where the second line follows from Markov's inequality. Thus, since ω_i maximizes the utility of agent i given ω_{-i} , we must have that

$$\frac{1}{2} \sum_{\mathbf{d}} \omega(\mathbf{d})[t_i(A_i) - 2\theta_i(A_i, \mathbf{d}_{-i})] \leq E_{\mathbf{d} \sim \omega}[t_i(X_i(\mathbf{d}))] = \sum_{\mathbf{d}} \omega(\mathbf{d})t_i(X_i(\mathbf{d})). \quad (3)$$

Combining (2) and (3), we conclude

$$\sum_{\mathbf{d}} \omega(\mathbf{d}) \sum_i t_i(X_i(\mathbf{d})) \geq \frac{1}{2c} \sum_i t_i(A_i) - \frac{1}{c} \sum_i \omega(\mathbf{d})t_i(X_i(\mathbf{d}))$$

which implies the desired inequality

$$E_{\mathbf{d} \sim \omega}[SW(\mathbf{X}(\mathbf{d}), \mathbf{t})] = \sum_{\mathbf{d}} \omega(\mathbf{d}) \sum_i t_i(X_i(\mathbf{d})) \geq \frac{1}{2(c+1)} \sum_i t_i(A_i) = \frac{1}{2(c+1)} SW_{opt}(\mathbf{t}).$$

□

We next show by way of an example that the analysis in Theorem 3.4 is tight.

Proposition 3.5. *For any $c \geq 2$, there is a combinatorial allocation problem \mathcal{P} such that the standard greedy algorithm \mathcal{A} provides a c -approximation for \mathcal{P} while the mixed Price of Anarchy for $\mathcal{M}_1(\mathcal{A})$ is $c + \Omega(\log c)$.*

Proof. Our problem is a combinatorial auction under two feasibility restrictions: first, no bidder can be allocated more than c objects. Second, certain agents can be allocated at most one object; say $N \subseteq [n]$ are these “singleton” agents. Let \mathcal{A} be the non-adaptive greedy algorithm with priority function $r(i, S, v) = v$ for $i \notin N$ and $r(i, S, v) = cv$ for $i \in N$. We note that this algorithm obtains a $(c+1)$ -approximation.

Consider the following instance of this problem. There are $2c^2$ objects, which we label a_{ij} and b_{ij} for $i, j \in [c]$. Let $\epsilon > 0$ be arbitrarily small. There are $4c$ agents, labelled A_i, B_i, B'_i , and C_i for $i \in [c]$. The singleton agents are the agents $\{C_i\}$. The types of the agents are as follows.

- For $i \in [c]$, agent A_i desires $\{a_{i1}, a_{i2}, \dots, a_{ic}\}$ for value c and $\{b_{i1}, b_{i2}, \dots, b_{ic}\}$ for value $1 + \epsilon$.
- For $i \in [c]$, agent B_i and B'_i both desire set $\{a_{i1}, a_{i2}, \dots, a_{i1}\}$ for value $(c - i)$.
- For $i \in [c]$, agent C_i desires $\{a_{i1}\}$ for value $1 - i/c$.

We can suppose that for any i , \mathcal{A} would break a tie between C_i, B_i , and B'_i in favour of C_i .

We now describe a mixed Nash equilibrium for this problem instance. Each agent A_i makes a single-minded bid of ϵ for set $\{b_{i1}, \dots, b_{ic}\}$. Each agent B_i and B'_i declares his valuation truthfully. Each agent C_i will declare his valuation truthfully with some probability p_i , and will otherwise declare the zero valuation. We choose $p_i = \frac{i}{i+1}$ for $i < c$, and $p_c = 1$.

We note that this distribution of declarations is indeed a Nash equilibrium. With probability 1, no agent B_i, B'_i , or C_i can obtain positive utility from any declaration, so their distributions over declarations that obtain 0 utility are necessarily optimal. Agent A_i obtains utility 1; his only hope for obtaining more utility is to declare a value less than $c - 1$ for set $\{a_{i1}, \dots, a_{ic}\}$. However, if he declares some value $c - Z$ with $Z > 1$, say with $X = \lceil Z \rceil$, then he can win the set only if bidders C_1, \dots, C_{X-1} all make single-minded bids, which occurs with probability $\frac{1}{2} \frac{2}{3} \dots \frac{X-1}{X} = \frac{1}{X} \leq \frac{1}{Z}$. Thus, for any Z , agent

A_i can obtain utility Z with probability at most $1/Z$, for an expected utility of at most 1. The given declaration by agent A_i is therefore optimal.

The optimal obtainable welfare in this example is c^2 , by allocating set $\{a_{i1}, \dots, a_{ic}\}$ to agent A_i for all i . In the equilibrium we've described, only objects $\{a_{11}, a_{12}, \dots, a_{1c}\}$ are allocated. For each i and $j < i$, object a_{1i} will be allocated to bidder B_j precisely if bidders C_1, \dots, C_{j-1} make single-minded bids but C_j does not, which occurs with probability $\frac{1}{j(j+1)}$. Object a_{1i} will be allocated to B_i, B'_i , or C_i with probability $\frac{1}{i}$. Noting that each of B_i, B'_i , and C_i has a per-item value of $1 - i/c$ for their desired sets, we conclude that the expected total value obtained is $\sum_{i \in [c]} (1 - \frac{1}{c} \sum_{j \leq i} \frac{1}{j}) = c - \frac{1}{c} \sum_{i \in [c]} H_i$, where H_i is the i th harmonic number. Since $H_i = \theta(\log i)$, we conclude that the expected social welfare is $c - \theta(\log c)$. The price of anarchy is therefore at least $\frac{c^2}{c - \theta(\log c)} = c + \theta(\log c)$, and hence is $c + \Omega(\log c)$. \square

3.3 Greedy Critical-Price Mechanisms

We have shown that the price of anarchy for $\mathcal{M}_1(\mathcal{A})$ is $c + \theta(\log c)$ when \mathcal{A} is a c -approximation. Intuitively, the extra logarithmic term is introduced because an agent may not know how to bid in order to obtain some set S and pay the minimum possible amount for it. This uncertainty is inherent in the first-price payment scheme. An alternative, the critical-price payment rule, does not exhibit this problem: under critical pricing, an agent that wins a set always pays the optimal price for it. Unfortunately, Lemma 3.2, which plays a crucial part in the proof of Theorem 3.4 fails to hold for the critical payment rule: there is no guarantee that agents will not overbid for mechanism $\mathcal{M}_{crit}(\mathcal{A})$. Indeed, there are settings in which an agent may be *strictly* better off by overbidding, as the example in section B shows.

We get around this issue of overbidding by imposing an assumption on the bidders that prevents them from doing so. We show that if the bidders are ex-post individually rational, and we apply the perturbed mechanism $\mathcal{M}_\mu(\mathcal{A})$ rather than $\mathcal{M}_{crit}(\mathcal{A})$, then we can assume that no bidders will overbid.

Claim 3.6. *If declaration profile \mathbf{d} for mechanism $\mathcal{M}_\mu(\mathcal{A})$ is made by ex-post individually-rational bidders, then it will satisfy $d_i(X_i(\mathbf{d})) \leq t_i(X_i(\mathbf{d}))$ for all i .*

Given that agents will not overbid, a simple modification of Theorem 3.4 yields a sharpened result.

Theorem 3.7. *Suppose \mathcal{A} is a c -approximate greedy allocation rule. Suppose also that the bidders are ex-post individually rational. Then for any $\mu > 0$, the expected welfare at any mixed Nash equilibrium of $\mathcal{M}_\mu(\mathcal{A})$ is a $(c + 1)(1 + \mu)$ approximation to the optimal welfare.*

The bound in Theorem 3.7 can be sharpened further to $c(1 + \mu)$ in certain important special cases. This follows directly from a corresponding improvement to Lemma 3.2, discussed in Section F.

4 Pure Nash Equilibria for Greedy Mechanisms

In this section we discuss pure equilibria in the full-information game for a mechanism with a greedy allocation rule. In such an equilibrium, agents declare according to a fixed profile \mathbf{d} , and no agent has incentive to deviate from d_i given that he knows \mathbf{d}_{-i} .

4.1 Pure Equilibria for First-Price Mechanisms

We first note that the approximation factor from Theorem 3.4 is improved when we restrict our attention to equilibria in pure strategies.

Theorem 4.1. *Suppose \mathcal{A} is a greedy c -approximate allocation rule for a combinatorial allocation problem. Then every pure Nash equilibrium of $\mathcal{M}_1(\mathcal{A})$ is a $(c + 1)$ -approximation.*

The power of Theorem 4.1 is marred by the fact that, for some problem instances, the mechanism $\mathcal{M}_1(\mathcal{A})$ is not guaranteed to have a pure Nash equilibrium. A simple example is given in Appendix D. Motivated by this example, we will consider a restricted set of allocation problems. This amounts to making an additional assumption on the space of feasible allocations. This assumption guarantees that the space is rich enough to allow agents to make conflicting bids.

Definition 4.2 (Blocking allocation problem). We say that an allocation problem is a *blocking allocation problem* if, for all $i, j \in [n]$, all partial allocations that do not include i or j , and all $S, T \subseteq M$, if S and T are feasible allocations to agent i , then there exists some feasible allocation $R \subseteq M$ to agent j such that no feasible allocation profile assigns R to agent j and either S or T to agent i .

For example, if we consider an allocation problem such that all allocated sets must be disjoint, then any auction that allows any player to obtain an arbitrary pair of objects (given that he bids highly enough on them) is a blocking allocation problem. We will additionally assume that the agents do not differentiate between outcomes whose utilities are within an additive difference of some arbitrarily small $\epsilon > 0$. Our solution concept will therefore be ϵ -Nash equilibrium. The price of anarchy bound from Theorem 4.1 also holds for ϵ -Nash equilibrium within an additive error that vanishes as ϵ tends to 0.

Finally, we make some assumptions on our greedy allocation rule. We will assume that the greedy algorithm is non-adaptive, and furthermore that it is *continuous*, meaning that its priority function $r(i, S, v)$ is a continuous function of v for all i and S . Many natural greedy algorithms are continuous; see Appendix J. We now show that $\mathcal{M}_1(\mathcal{A})$ has a pure ϵ -Nash equilibrium for any blocking allocation problem, when \mathcal{A} is non-adaptive and continuous.

Theorem 4.3. *Suppose \mathcal{A} is a c -approximate non-adaptive continuous greedy allocation rule for a blocking allocation problem. Then for any $\epsilon > 0$, $\mathcal{M}_1(\mathcal{A})$ has a pure ϵ -Nash equilibrium.*

4.2 Pure Equilibria for Critical-Price Mechanisms

Theorem 4.3 requires that the underlying allocation problem is blocking. In this section we consider an alternative pure equilibrium construction for non-blocking allocation problems. For this we will use the critical-price mechanism. The critical-price mechanism always has a pure equilibrium for any greedy algorithm and any combinatorial allocation problem: given any feasible allocation profile \mathbf{A} , the declaration profile in which each agent i bids very highly and single-mindedly on set A_i is an equilibrium. However, this mechanism has unbounded price of anarchy due to this ease of constructing equilibria: any allocation, no matter how bad an approximation to the optimal social welfare, could occur at equilibrium.

Our goal is to guarantee the existence of equilibria like those described above, but also retain the price of anarchy result from Theorem 4.1. To do this, we require that agents do not overbid. Recall that in Theorem 3.7 we prevented overbidding by assuming agents were ex-post individually rational, and modifying the payment scheme to allow a small probability that agents would have to pay according to their declared values. That approach will not work here, as the pure equilibria for critical-price mechanisms described above will not necessarily be equilibria under this modified payment scheme.

Rather than use the ex-post individual rationality assumption and mechanism $\mathcal{M}_\mu(\mathcal{A})$, we use mechanism $\mathcal{M}_{crit}(\mathcal{A})$ and directly assume that agents do not overbid. We motivate this assumption by noting that, when \mathbf{d}_{-i} is known, agent i never gains benefit from overbidding. The assumption that \mathbf{d}_{-i} is fixed, motivated by our pure equilibrium setting, is necessary for this result: as shown in Section 3, an agent may derive benefit from overbidding when the declarations of other agents are not fixed.

Claim 4.4. *For each i and \mathbf{d}_{-i} , there is a utility-maximizing declaration d_i for mechanism $\mathcal{M}_{crit}(\mathcal{A})$ such that $d_i(S) \leq t_i(S)$ for all $S \subseteq M$.*

The previous claim shows that the set of strategies in which an agent does not overbid is weakly dominant for that agent. We now show that if \mathcal{A} is a standard greedy algorithm, then there is a pure Nash equilibrium for $\mathcal{M}_{crit}(\mathcal{A})$ using strategies from this dominant set. Moreover, if agents only make declarations from this dominant set, then we obtain the same price of anarchy bound as in Theorem 4.1.

Theorem 4.5. *Suppose \mathcal{A} is a c -approximate standard greedy allocation rule. Then $\mathcal{M}_{crit}(\mathcal{A})$ has a pure Nash equilibrium in which no agent overbids, and if players do not overbid then the price of anarchy for $\mathcal{M}_{crit}(\mathcal{A})$ in pure strategies is $(c + 1)$.*

5 Combining Mechanisms

A common technique in the design of allocation rules is to consider both a greedy rule that favours allocation of small sets, and a simple rule that allocates all objects to a single bidder, and apply whichever

solution obtains the better result [21, 7, 4]. When bidders are single-minded, such a combination rule will be incentive-compatible [21]. We would like to extend our results to cover rules of this form, but the price of anarchy for such a rule (with either the first-price or critical-price payment scheme) may be much worse than its combinatorial approximation ratio (see Section I). We therefore consider a different way to combine two rules: we implement each rule as a separate mechanism, then randomly choose between the two mechanisms with equal probability. For many examples of interest (see Section J) the resulting randomized allocation rule obtains the same worst-case combinatorial approximation ratio as applying the better of the two rules for each input. Moreover, the price of anarchy results of this paper can be made to carry over to such randomized mechanisms, as we now formalize.

Let \mathcal{A} be any greedy allocation rule that never allocates M to any agent, and let \mathcal{A}' be the allocation rule that allocates M to the agent i that maximizes $d_i(M)$. The restriction on \mathcal{A} is motivated by our intuition that \mathcal{A} favours allocations of small sets; it is without loss for many algorithms of interest (see Section J). We write $\mathcal{M}_1(\mathcal{A}, \mathcal{A}')$ for the mechanism that flips a fair coin, and if it lands heads it executes $\mathcal{M}_1(\mathcal{A})$, otherwise executes $\mathcal{M}_1(\mathcal{A}')$. We define $\mathcal{M}_{crit}(\mathcal{A}, \mathcal{A}')$ and $\mathcal{M}_\mu(\mathcal{A}, \mathcal{A}')$ similarly. For these mechanisms, we will *allow input valuations to be non-monotone* with respect to set M ; that is, we allow declarations in which $d_i(M) < d_i(S)$ for $S \subseteq M$. Note then that our mechanism is not technically a direct revelation mechanism, as an agent's input is not necessarily a valid valuation function.

Theorem 5.1. *Suppose that $\mathcal{A}, \mathcal{A}'$ are as described above, and for every declaration profile \mathbf{d} , $SW(\mathcal{A}, \mathbf{d}) + SW(\mathcal{A}', \mathbf{d}) \geq \frac{1}{c} SW_{opt}(\mathbf{d})$. Then:*

- $\mathcal{M}_1(\mathcal{A}, \mathcal{A}')$ obtains a $2(c + O(\log c))$ approximation at every mixed equilibrium,
- $\mathcal{M}_\mu(\mathcal{A}, \mathcal{A}')$ obtains a $2(c + 1)$ approximation at every mixed equilibrium, assuming ex-post individually rational bidders,
- $\mathcal{M}_1(\mathcal{A}, \mathcal{A}')$ obtains a $2(c + 1)$ approximation at every pure equilibrium, and has a pure equilibrium if \mathcal{A} is continuous and non-adaptive and the allocation problem is blocking.
- $\mathcal{M}_{crit}(\mathcal{A}, \mathcal{A}')$ obtains a $2(c + 1)$ approximation at every pure equilibrium, and has a pure equilibrium, if agents do not overbid and \mathcal{A} is a standard greedy algorithm.

6 Conclusions and Open Problems

We have demonstrated that many greedy algorithms for combinatorial allocation problems can be implemented as deterministic mechanisms without much loss to their approximation ratios at any mixed Bayesian Nash equilibrium. This has a number of applications, such as a combinatorial auction with $O(\sqrt{m})$ price of anarchy, an efficient mechanism for unit-job profit-maximizing scheduling with price of anarchy 3, a mechanism for the (multiple demand) unsplittable flow problem with price of anarchy $O(m^{\frac{1}{B-1}})$, and others. A more detailed list of applications is deferred to Appendix J. We extended this analysis to pure equilibria for rich subclasses of problems and algorithms.

There are a number of immediate questions left open in our results. The first is to improve the price of anarchy bound in Theorem 3.7; it would be interesting to determine whether any greedy allocation rule can be implemented deterministically so that there is no loss in approximation ratio at any mixed equilibrium, without assumptions on the agents. Another is to relax the assumptions of Theorems 4.3 and 4.5, finding other classes of problems for which pure equilibria are guaranteed to exist.

More generally, the price of anarchy solution concept can be applied to other mechanism design problems and other types of algorithms. We ask: given a mechanism design problem, when can an algorithm for the underlying optimization problem be converted into a mechanism that obtains (nearly) the same approximation ratio at every BNE? Even a partial resolution would be an important step in understanding the relationship between computational issues and Bayesian Nash implementability.

Acknowledgements

We are thankful to Jason Hartline, Nicole Immorlica, and Silvio Micali for helpful discussions.

References

- [1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*, 2001.
- [2] 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.
- [3] P. Baptiste. Polynomial time algorithms for minimizing the weighted number of late jobs on a single machine with equal processing times. *Journal of Scheduling*, 2:245–252, 1999.
- [4] 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.
- [5] A. Borodin and B. Lucier. Greedy mechanism design for truthful combinatorial auctions. Working Paper, 2009.
- [6] A. Borodin, M. N. Nielsen, and C. Rackoff. (incremental) priority algorithms. In *Proc. 13th ACM Symp. on Discrete Algorithms*, 2002.
- [7] P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanism design. In *Proc. 36th ACM Symp. on Theory of Computing*, 2005.
- [8] G. Christodoulou, A. Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *Proc. 35th Intl. Colloq. on Automata, Languages and Programming*, pages 820–832, 2008.
- [9] S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *Proc. 9th ACM Conf. on Electronic Commerce*, 2007.
- [10] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proc. 37th ACM Symp. on Theory of Computing*, 2006.
- [11] 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.
- [12] 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.
- [13] J. Hastad. Some optimal inapproximability results. In *Proc. 29th ACM Symp. on Theory of Computing*, 1997.
- [14] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.
- [15] N. Immorlica, D. Karger, E. Nikolova, and R. Sami. First-price path auctions. In *Proc. 7th ACM Conf. on Electronic Commerce*, 2005.
- [16] M. Jackson. A crash course in implementation theory. In *Social Choice and Welfare*, 2001.
- [17] 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.
- [18] B. Lehmann, D. Lehmann, and N. Nisan. Truthful and near-optimal mechanism design via linear programming. In *Proc. 3rd ACM Conf. on Electronic Commerce*, 2001.
- [19] 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.
- [20] J. Mestre. Greedy in approximation algorithms. In *Proc. 14th European Symp. on Algorithms*, 2006.
- [21] A. Mu’alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64:612–631, 2008.
- [22] N. Nisan. The communication complexity of approximate set packing and covering. In *Proc. 29th Intl. Colloq. on Automata, Languages and Programming*, 2002.

- [23] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. 31st ACM Symp. on Theory of Computing*, pages 129–140. ACM Press, 1999.
- [24] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [25] C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, 2008.
- [26] Hal Varian. Position auctions. In *Working Paper, UC Berkeley*, 2006.

A Bayesian Price of Anarchy for First-Price Auctions

In a Bayesian setting, we suppose that the true types of the agents are drawn from a known probability distribution D over the set of valuation profiles. We assume that $D = D_1 \times \dots \times D_n$ is the product of independent distributions, where $D_i(t_i)$ is the probability that agent i has type t_i . We write $SW_{opt}(D)$ for the expectation of $SW_{opt}(\mathbf{t})$, where \mathbf{t} is chosen according to distribution D .

Given a type t_i for agent i , let $\omega_i^{t_i}$ be a distribution of declarations for agent i , parameterized by t_i . We think of $\omega_i^{t_i}$ as the bidding strategy employed by agent i given that his true type is t_i . We write $\omega^{\mathbf{t}} = \omega_1^{t_1} \times \dots \times \omega_n^{t_n}$. The set of distributions $\{\omega^{\mathbf{t}}\}$, over all choices of \mathbf{t} , forms a *Bayesian Nash Equilibrium* if for every $i \in [n]$, and every t_i in the support of D_i , agent i maximizes his expected utility by making a declaration drawn from distribution $\omega_i^{t_i}$. Here the expectation is with respect to \mathbf{t}_{-i} being drawn from D_{-i} , and the declarations of the other players then being drawn from $\omega_{-i}^{\mathbf{t}_{-i}}$. In other words, assuming types are drawn from D and players follow the strategies set out in $\omega^{\mathbf{t}}$, no player has incentive to unilaterally deviate. We write $SW(D, \omega^{\mathbf{t}})$ for the expected social welfare given type distribution D and strategies profiles $\omega^{\mathbf{t}}$.

The (*mixed*) *Bayesian price of anarchy* is

$$\max_{D, \{\omega^{\mathbf{t}}\}} \frac{SW_{opt}(D)}{SW(D, \{\omega^{\mathbf{t}}\})}$$

where the maximum is taken over all type distributions D and Bayesian Nash equilibria $\{\omega^{\mathbf{t}}\}$ for D .

We complete the proof of Theorem 3.4 by showing that the price of anarchy of $\mathcal{M}_1(\mathcal{A})$ is $c + O(\log c)$, and extending the argument to Bayesian settings.

Fix a distribution $D = D_1 \times \dots \times D_n$ over type profiles. Let $\{\omega^{\mathbf{t}}\}$ be a Bayesian Nash equilibrium with respect to D .

Choose some \mathbf{t} in the support of D , and let \mathbf{d} be a declaration profile in the support of $\omega^{\mathbf{t}}$. Let $\mathbf{A}^{\mathbf{t}} = A_1^{\mathbf{t}}, \dots, A_n^{\mathbf{t}}$ denote an optimal allocation for \mathbf{t} . By Lemma 3.2 and Lemma 3.3,

$$\sum_{i \in [n]} t_i(X_i(\mathbf{d})) \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i}).$$

Summing over all choices of \mathbf{d} with respect to $\omega^{\mathbf{t}}$, we have

$$\sum_{\mathbf{d}} \omega^{\mathbf{t}}(\mathbf{d}) \sum_i t_i(X_i(\mathbf{d})) \geq \sum_{\mathbf{d}} \omega^{\mathbf{t}}(\mathbf{d}) (1/c) \sum_i \theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i}) \quad (4)$$

Choose some $i \in [n]$. We now present a parameterized bound on $E_{\mathbf{d}}[\theta_i(S, \mathbf{d}_{-i})]$ with respect to $t_i(S)$ and $t_i(X_i(\mathbf{d}))$, for any set S .

Claim A.1. *For any set S and any $k > 1$, if d_i is in the support of $\omega_i^{t_i}$, then*

$$E_{\mathbf{v}_{-i}, \mathbf{d}_{-i}}[\theta_i(S, \mathbf{d}_{-i})] \geq \left(1 - \frac{1 + \log k}{k}\right) t_i(S) - (1 + \log k) E_{\mathbf{v}_{-i}, \mathbf{d}_{-i}}[t_i(X_i(\mathbf{d}))]$$

where expectations are over \mathbf{v}_{-i} drawn from D_{-i} and \mathbf{d}_{-i} drawn from $\omega_{-i}^{\mathbf{v}_{-i}}$.

Proof. For brevity, we will write $E[\theta_i] = E_{\mathbf{v}_{-i}, \mathbf{d}_{-i}}[\theta_i(S, \mathbf{d}_{-i})]$ and $E[t_i] = E_{\mathbf{v}_{-i}, \mathbf{d}_{-i}}[t_i(X_i(\mathbf{d}))]$, where each expectation is over $\mathbf{v}_{-i} \sim D_{-i}$ and $\mathbf{d}_{-i} \sim \omega_{-i}^{\mathbf{v}_{-i}}$.

First note that if $t_i(S) \leq E[t_i]$ the result is trivially true (since $1 + \log k > 1 - \frac{1 + \log k}{k}$ and $E[\theta_i] \geq 0$), so suppose otherwise. If $E[t_i] = 0$, it must be that $\theta_i(S, \mathbf{d}_{-i}) \geq t_i(S)$ for all $\mathbf{d}_{-i} \sim \omega_{-i}^{\mathbf{v}_{-i}}$, since otherwise agent i could improve his utility by bidding slightly less than $t_i(S)$ for set S . This would imply the desired result, so we can assume $E[t_i] > 0$.

Let $r = t_i(S)/E[t_i]$; by the previous paragraph we have $r \in [1, \infty)$. Choose any $Z \in [1, r]$. Suppose for contradiction that $Pr_{\mathbf{d}_{-i}}[\theta_i(S, \mathbf{d}_{-i}) < t_i(S) - ZE[t_i]] > 1/Z$. Recalling that the expected utility of agent i on declaration \mathbf{d} is at most $E[t_i]$, we see that agent i could improve his utility by making a single-minded bid for set S at value $t_i(S) - ZE[t_i]$ with probability 1. This is a contradiction, so we must conclude $Pr[\theta_i(S, \mathbf{d}_{-i}) < t_i(S) - ZE[t_i]] \leq 1/Z$ for all $Z \in [1, r]$. Since we also know $\theta_i(S, \mathbf{d}_{-i}) \geq 0$ with probability 1, we conclude $E[\theta_i] \geq t_i(S) - (1 + \log r)E[t_i]$. We now proceed by cases.

Case 1: $r \leq k$: then $E[\theta_i] \geq t_i(S) - (1 + \log r)E[t_i] \geq t_i(S) - (1 + \log k)E[t_i]$.

Case 2: $r \geq k$: then $E[\theta_i] \geq t_i(S) - (1 + \log r)E[t_i] = (1 - \frac{1 + \log r}{r})t_i(S) \geq (1 - \frac{1 + \log k}{k})t_i(S)$.

Hence, for all $k > 1$, either $E[\theta_i] \geq t_i(S) - (1 + \log k)E[t_i]$ or $E[\theta_i] \geq (1 - \frac{1 + \log k}{k})t_i(S)$, which implies the desired result. \square

Applying this Claim A.1 to $A_i^{\mathbf{t}}$, we conclude that for all $i \in [n]$, \mathbf{t} , and $k > 1$,

$$\begin{aligned} \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}_{-i}} \omega_{-i}^{\mathbf{v}}(\mathbf{d}_{-i}) \theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i}) &\geq (1 - \frac{1 + \log k}{k}) t_i(A_i^{\mathbf{t}}) \\ &\quad - (1 + \log k) \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}_{-i}} \omega_{-i}^{\mathbf{v}}(\mathbf{d}_{-i}) t_i(X_i(\mathbf{d})) \end{aligned} \quad (5)$$

Summing over i and \mathbf{t} , we have

$$\begin{aligned} \sum_{\mathbf{t}} D(\mathbf{t}) \sum_{i \in [n]} \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}_{-i}} \omega_{-i}^{\mathbf{v}}(\mathbf{d}_{-i}) \theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i}) \\ \geq (1 - \frac{1 + \log k}{k}) \sum_{\mathbf{t}} D(\mathbf{t}) \sum_{i \in [n]} t_i(A_i^{\mathbf{t}}) \\ - (1 + \log k) \sum_{\mathbf{t}} D(\mathbf{t}) \sum_{i \in [n]} \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}_{-i}} \omega_{-i}^{\mathbf{v}}(\mathbf{d}_{-i}) t_i(X_i(\mathbf{d})) \end{aligned} \quad (6)$$

Note $\sum_{\mathbf{t}} D(\mathbf{t}) \sum_{i \in [n]} t_i(A_i^{\mathbf{t}})$ is precisely $SW_{opt}(D)$. Also, we note

$$\begin{aligned} \sum_{\mathbf{t}} D(\mathbf{t}) \sum_{i \in [n]} \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}_{-i}} \omega_{-i}^{\mathbf{v}}(\mathbf{d}_{-i}) \theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i}) &= \sum_{\mathbf{t}} D(\mathbf{t}) \sum_{i \in [n]} \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}} \omega^{\mathbf{v}}(\mathbf{d}) \theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i}) \\ &= \sum_{\mathbf{t}} D(\mathbf{t}) \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}} \omega^{\mathbf{v}}(\mathbf{d}) \sum_{i \in [n]} \theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i}) \\ &\leq \sum_{\mathbf{t}} D(\mathbf{t}) \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}} \omega^{\mathbf{v}}(\mathbf{d}) c \sum_{i \in [n]} d_i(X_i(\mathbf{d})) \\ &= \sum_{\mathbf{v}} D(\mathbf{v}) \sum_{\mathbf{d}} \omega^{\mathbf{v}}(\mathbf{d}) c \sum_{i \in [n]} d_i(X_i(\mathbf{d})) \\ &= cSW(D, \{\omega^{\mathbf{t}}\}) \end{aligned}$$

So we conclude

$$cSW(D, \{\omega^{\mathbf{t}}\}) \leq (1 - \frac{1 + \log k}{k}) SW_{opt}(D) - (1 + \log k) SW(D, \{\omega^{\mathbf{t}}\})$$

Setting $k = c$ and rearranging yields

$$SW(D, \{\omega^{\mathbf{t}}\}) \geq \left(\frac{1}{c+1+\log c}\right) \left(\frac{c}{c-1-\log c}\right) SW_{opt}(D) = \frac{1}{c+O(\log c)} SW_{opt}(D),$$

for a price of anarchy of $c + O(\log c)$. This completes the proof of Theorem 3.4.

B Overbidding example from section 3.3

Example B.1. Consider a combinatorial auction with 3 objects, $\{a, b, c\}$, and 3 bidders, under the additional restriction that each agent can be allocated at most one object. Let \mathcal{A} be the standard greedy algorithm for this problem. Suppose the types of the players are as follows: $t_1(b) = 2$, $t_1(c) = 4$, $t_2(c) = 3$, $t_3(a) = 1$, $t_3(b) = 6$, and all other values are 0. Consider the following bidding strategies for agents 2 and 3: bidder 2 declares truthfully with probability 1, and bidder 3 either declares single-mindedly for a with value 1, or single-mindedly for b with value 6, each with equal probability.

How should agent 1 declare to maximize utility? We can limit our attention to pure strategies, by linearity of expectation. Suppose agent 1 does not overbid and declares at most 2 for object b . If he also declares at least 3 for object c , then he wins c with probability 1 for an expected utility of 1. If he doesn't declare at least 3 for object c , then he wins b with probability 1/2 and nothing otherwise, again for an expected utility of 1. So agent 1 can gain a utility of at most 1 if he does not overbid. If, however, he declares 5 for b and 4 for c , then he wins b with probability 1/2 and wins c otherwise, for an expected utility of 3/2. The reader can verify that if agent 1 bids in this way, the resulting combination of strategies forms a mixed Nash equilibrium. Thus, in mixed equilibria, an agent may strictly improve his utility by overbidding.

C Other proofs omitted in section 3

Proof of Lemma 3.2: Let \mathbf{d}' be the declaration profile in which d_i' is the single-minded declaration for set $X_i^{\mathcal{A}}(\mathbf{d})$ at value $d_i(X_i^{\mathcal{A}}(\mathbf{d}))$ for all i . Repeated application of Lemma 3.1 implies that allocations and critical prices for all agents are the same for inputs \mathbf{d} and \mathbf{d}' .

Choose any $\epsilon > 0$. Let d_i'' be the single-minded declaration for set A_i at value $\theta_i(A_i, \mathbf{d}_{-i}) - \epsilon$. Let d_i^* be the maximum of d_i' and d_i'' . Repeated application of Lemma 3.1 again implies that allocations and critical prices of \mathbf{d} and \mathbf{d}' are identical, using the fact that A_i cannot be allocated to agent i unless $A_i \supseteq X_i^{\mathcal{A}}(\mathbf{d})$.

Since \mathcal{A} is a c -approximation, we conclude that $SW(\mathbf{X}^{\mathcal{A}}(\mathbf{d}^*), \mathbf{d}^*) \geq \frac{1}{c} SW(\mathbf{A}, \mathbf{d}^*) - n\epsilon$. Since $\mathbf{X}^{\mathcal{A}}(\mathbf{d}^*) = \mathbf{X}^{\mathcal{A}}(\mathbf{d})$, the result follows by taking the limit as $\epsilon \rightarrow 0$. \square

Proof of Claim 3.6: If $d_i(S) > t_i(S)$, where $S = X_i(\mathbf{d})$, then with positive probability agent i will be allocated set S and pay $d_i(S)$, for a utility of $t_i(S) - d_i(S) < 0$. Since agent i could always choose to bid so that he obtains non-negative utility with probability 1, this contradicts the ex-post individual rationality of agent i . \square

Proof of Lemma 3.3: Suppose for contradiction that $d_i(S) > t_i(S)$. Define d_i' by $d_i'(S) = \min\{d_i(S), t_i(S)\}$ for all $S \subseteq M$. Note that d_i' satisfies monotonicity and is therefore a valid valuation function. Recall that $u_i(\mathbf{d})$ denotes the utility obtained by agent i given declaration profile \mathbf{d} ; since ω forms a mixed NE, choosing a declaration according to distribution ω_i must maximize the expected utility of agent i . Since d_i has positive probability in ω_i , linearity of expectation implies that $E_{\mathbf{d}_{-i} \sim \omega_{-i}}[u_i(d_i, \mathbf{d}_{-i})] \geq E_{\mathbf{d}_{-i} \sim \omega_{-i}}[u_i(d_i', \mathbf{d}_{-i})]$.

Fix any \mathbf{d}_{-i} and let $S_i = X_i(d_i', \mathbf{d}_{-i})$, $T_i = X_i(d_i, \mathbf{d}_{-i})$. If $d_i(T_i) \leq t_i(T_i)$, then $u_i(d_i, \mathbf{d}_{-i}) \leq 0 \leq u_i(d_i', \mathbf{d}_{-i})$. If, on the other hand, $d_i(T_i) > t_i(T_i)$, then we claim $S_i = T_i$. Otherwise, since critical prices are unchanged, it must be that on some iteration of \mathcal{A} , $r(i, S_i, d_i(S_i)) \geq r(i, T_i, d_i(T_i))$ and $r(i, S_i, d_i'(T_i)) \geq r(i, T_i, d_i'(S_i))$, but this leads to a contradiction since $d_i'(T_i) = d_i(T_i)$ and $d_i'(S_i) \leq d_i(S_i)$, and r is monotone. Thus $S_i = T_i$, and hence $u_i(d_i, \mathbf{d}_{-i}) = u_i(d_i', \mathbf{d}_{-i})$.

We conclude $u_i(d_i', \mathbf{d}_{-i}) \geq u_i(d_i, \mathbf{d}_{-i})$, and this inequality is strict when $X_i(d_i, \mathbf{d}_{-i}) = S$. Since this event occurs with positive probability, we conclude $E_{\mathbf{d} \sim \omega}[u_i(d_i, \mathbf{d}_{-i})] < E_{\mathbf{d}_{-i} \sim \omega_{-i}}[u_i(d_i', \mathbf{d}_{-i})]$, a contradiction. \square

Proof of Theorem 3.7: Fix a true type profile \mathbf{t} . Let $\omega = (\omega_1, \dots, \omega_n)$ be a probability distribution on the set of all possible declarations, and suppose that ω forms a mixed Nash equilibrium. Let A_1, \dots, A_n denote an optimal allocation for \mathbf{t} . Following Theorem 3.4, we have that

$$\begin{aligned} \sum_i v_i(X_i(\mathbf{d})) &\geq (1/c) \sum_i \theta_i(A_i, \mathbf{d}_{-i}) \\ &= (1/c) \sum_i t_i(A_i) - (1/c) \sum_i [t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})]. \end{aligned}$$

Summing over all declaration profiles, with respect to probability function ω , we have that

$$\begin{aligned} \sum_{\mathbf{d}} \omega(\mathbf{d}) \sum_i t_i(X_i(\mathbf{d})) &\geq \sum_{\mathbf{d}} \omega(\mathbf{d}) (1/c) \sum_i t_i(A_i) - \sum_{\mathbf{d}} \omega(\mathbf{d}) (1/c) \sum_i [t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})] \\ &= (1/c) \sum_i t_i(A_i) - (1/c) \sum_i \omega(\mathbf{d}) [t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})]. \end{aligned}$$

Let d_i'' be the single-minded declaration for set A_i at value $t_i(A_i)$. Let us consider the expected utility for player i if he were to declare d_i'' with probability 1. We have

$$\begin{aligned} E_{\mathbf{d}_{-i} \sim \omega_{-i}}[t_i(X_i(d_i'', \mathbf{d}_{-i}))] &= \sum_{\mathbf{d}_{-i}} \omega_{-i}(\mathbf{d}_{-i}) \Pr[t_i(A_i) \geq \theta(A_i, \mathbf{d}_{-i})] (t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})) \\ &\geq \sum_{\mathbf{d}_{-i}} \omega_{-i}(\mathbf{d}_{-i}) (t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})) \\ &= \sum_{\mathbf{d}} \omega(\mathbf{d}) (t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})) \end{aligned}$$

Thus, since ω_i maximizes the utility of agent i given ω_{-i} , we must have that

$$\omega(\mathbf{d}) [t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})] \leq \omega(\mathbf{d}) t_i(X_i(\mathbf{d})).$$

We conclude

$$\sum_{\mathbf{d}} \omega(\mathbf{d}) \sum_i t_i(X_i(\mathbf{d})) \geq (1/c) \sum_i t_i(A_i) - (1/c) \sum_i \omega(\mathbf{d}) t_i(X_i(\mathbf{d}))$$

which implies

$$\sum_{\mathbf{d}} \omega(\mathbf{d}) \sum_i t_i(X_i(\mathbf{d})) \geq (1/(c+1)) \sum_i t_i(A_i)$$

as required. \square

D Example omitted from section 4

We provide an example of a combinatorial allocation problem, greedy allocation rule \mathcal{A} , and type profile \mathbf{t} such that $\mathcal{M}_1(\mathcal{A})$ has no pure equilibria when agents have types \mathbf{t} .

Example D.1. Consider an instance of the combinatorial auction problem where each agent can be assigned at most one object, and let \mathcal{A} be the standard greedy algorithm. Suppose we have two objects, $M = \{a, b\}$, and three agents. Suppose the true types of the agents are as follows: $t_1(a) = 4$, $t_1(b) = 2$, $t_2(a) = 3$, $t_2(b) = 0$, $t_3(a) = 0$, and $t_3(b) = 3$.

We now prove that no pure ϵ -Nash equilibrium exists for this example, for any $\epsilon \in (0, 1)$. Assume for contradiction that there is a pure ϵ -Nash equilibrium \mathbf{d} for \mathbf{t} and mechanism $\mathcal{M}_1(\mathcal{A})$.

We know that $d_1(\{b\}) \leq 2$ and $d_2(\{b\}) \leq 0$ (by Lemma 3.3). Thus it must be that $X_3(\mathbf{d}) = \{b\}$, since otherwise agent 3 could change his declaration to win $\{b\}$ and increase his utility. Thus, since

agent 1 does not win item $\{b\}$, we conclude that $X_1(\mathbf{d}) = \{a\}$, since otherwise agent 1 could change his declaration to win $\{a\}$ and increase his utility.

Now note that if $d_1(\{a\}) < 3$, agent 2 could increase his utility by making a winning declaration for $\{a\}$. Thus $d_1(\{a\}) \geq 3$, and hence $u_1(\mathbf{d}) \leq 4 - 3 = 1$. This also implies that $d_1(\{a\}) > d_1(\{b\})$, so agent 3 would win $\{b\}$ regardless of his bid. Thus, to maximize his utility, it must be that $d_3(\{b\}) \leq \epsilon$. But then agent 1 could improve his utility by changing his declaration and bidding 0 for $\{a\}$ and 2ϵ for $\{b\}$, obtaining utility $2 - 2\epsilon > 1$. Therefore \mathbf{d} is not an equilibrium, a contradiction.

E Proofs omitted in section 4

Proof of Theorem 4.1: Fix type profile \mathbf{t} and suppose \mathbf{d} is a pure Nash equilibrium. Let A_1, \dots, A_n be an optimal allocation. Since \mathcal{A} is a c -approximate allocation rule, we know (following Theorem 3.4)

$$\begin{aligned} \sum_i t_i(X_i(\mathbf{d})) &\geq \sum_i d_i(X_i(\mathbf{d})) \\ &\geq (1/c) \sum_i \theta_i(A_i, \mathbf{d}_{-i}) \\ &= (1/c) \sum_i t_i(A_i) - (1/c) \sum_i [t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})]. \end{aligned}$$

Let d_i'' be the single-minded declaration for set A_i at value $\theta_i(A_i) + \epsilon$. The expected utility for player i if he were to bid d_i'' is $t_i(A_i) - \theta_i(A_i)$, so since \mathbf{d} forms an equilibrium we conclude $t_i(X_i(\mathbf{d})) \geq t_i(A_i) - \theta_i(A_i)$. This implies

$$\sum_i t_i(X_i(\mathbf{d})) \geq (1/c) \sum_i t_i(A_i) - (1/c) \sum_i t_i(X_i(\mathbf{d}))$$

which implies

$$\sum_i t_i(X_i(\mathbf{d})) \geq (1/(c+1)) \sum_i t_i(A_i)$$

as required. \square

Proof of Theorem 4.3: We will construct an equilibrium \mathbf{d} explicitly. The intuition behind the construction is as follows. Each agent i will bid truthfully for set $X_i(\mathbf{t})$. Other agents will then place “blocking bids” that conflict with both $X_i(\mathbf{t})$ and any other set that agent i might desire, with rank slightly below the rank of the bid for $X_i(\mathbf{t})$. These blocking bids are not allocated by the algorithm, since they conflict with the bid for $X_i(\mathbf{t})$, but they guarantee that agent i cannot increase his utility by abandoning set $X_i(\mathbf{t})$ and attempting to obtain a different set.

More formally, we will define declaration \mathbf{d} incrementally as follows. First, initialize d_i be the single-minded valuation of $t_i(X_i(\mathbf{t}))$ for set $X_i(\mathbf{t})$, for each $i \in [n]$. Define $r_i = r(i, X_i(\mathbf{t}), t_i(X_i(\mathbf{t})))$ for each agent i . Then, for every pair of agents i, j such that $r_i > r_j$, and every set S such that $X_i(\mathbf{t}) \not\subseteq S$ and $t_i(S) \geq \theta_i(S, \mathbf{d}_{-i})$, find a set R such that R is a feasible allocation to agent j given the allocations to agents k with $r_k > r_i$, and additionally no outcome allocates R to bidder j and either S or $X_i(\mathbf{t})$ to bidder i . Such an R must exist from the definition of a blocking allocation problem, and moreover $X_j(\mathbf{t}) \not\subseteq R$ (since there *does* exist a feasible allocation that allocates $X_j(\mathbf{t})$ to agent j and $X_i(\mathbf{t})$ to agent i , namely $\mathbf{X}(\mathbf{t})!$). Set $d_j(R)$ so that $r(i, S, t_i(S) - \epsilon) \leq r(j, R, d_j(R)) \leq r(i, S, t_i(S) - \epsilon/2)$, increasing declarations for other sets if necessary to retain monotonicity. Such a value of $d_j(R)$ must exist, since \mathcal{A} is a continuous algorithm with bounded approximation ratio. Note then that $r(j, R, d_j(R)) < r_i$, so this declaration for set R will fall below its critical price. Finally, for the player i with minimal r_i , adjust d_i by setting $d_i(X_i(\mathbf{t})) = \epsilon$.

We have finished the definition of \mathbf{d} . Note that $\mathbf{X}(\mathbf{d}) = \mathbf{X}(\mathbf{t})$, as the adjustments to \mathbf{d} only added declared values that fell below critical prices. For any agent i , no change in declaration can improve his utility by more than ϵ , since for any set S , $t_i(S) \leq \theta_i(S, \mathbf{d}_{-i}) + \epsilon$. Hence \mathbf{d} is a pure Nash equilibrium. \square

Proof of Claim 4.4: Choose any utility-maximizing d_i' , say with $X_i(d_i', \mathbf{d}_{-i}) = T$. Let d_i be the single-minded declaration for set T of value $t_i(T)$. Then, by the monotonicity of t_i , we know $d_i(S) \leq t_i(S)$ for all S . Since $X_i(d_i, \mathbf{d}_{-i}) = T$ we have $u_i(d_i', \mathbf{d}_{-i}) = u_i(d_i, \mathbf{d}_{-i})$. Thus, since d_i' is utility-maximizing, d_i is as well. \square

Proof of Theorem 4.5: Define \mathbf{d} by having d_i be the single-minded bid for set $X_i(\mathbf{t})$ at value $t_i(X_i(\mathbf{t}))$. Then the critical price for each agent's winning set is 0, so $u_i(\mathbf{d}) = t_i(X_i(\mathbf{t}))$. Since \mathcal{A} ranks allocations by value, no agent can obtain a higher utility by winning any other set, so this is an equilibrium.

The price of anarchy result follows precisely the proof of the corresponding results for $\mathcal{M}_\mu(\mathcal{A})$, using the assumption that no agent will overbid. \square

F Tightening Results for Special Cases

In this section we show how to tighten the results of Lemma 3.2 for certain special cases of allocation problems and greedy algorithms. This allows us to obtain sharper bounds in Theorems 3.7 and 4.1. We say that a combinatorial allocation problem is *player symmetric* if the feasibility constraints do not depend on the labelling of the players, and *object symmetric* if they do not depend on the labelling of the objects. We say that a greedy algorithm is *player symmetric* if its ranking function does not depend on its first parameter, and we say that it is *object symmetric* if its ranking function depends only on the cardinality of the second parameter.

Lemma F.1. *If \mathcal{A} is a player-symmetric greedy algorithm and a $c(n)$ -approximation whenever all declarations are single-minded, then for any declaration profile \mathbf{d} and allocation profile $\mathbf{A} = A_1, \dots, A_n$, $\sum_{i \in [n]} d_i(X_i(\mathbf{d})) \geq \frac{1}{c(2n)} \sum_{i \in [n]} \theta_i(A_i, \mathbf{d}_{-i})$*

Proof. We define \mathbf{d}' as in Lemma 3.2. We then define \mathbf{d}'' by adding n additional bidders, $1', \dots, n'$, where $d_{i'}''$ is the single-minded declaration for set A_i at value $\theta_i(A_i, \mathbf{d}_{-i}) - \epsilon$. Player symmetry implies that $\mathbf{X}(\mathbf{d}') = \mathbf{X}(\mathbf{d}'')$ (meaning that each additional player is allocated \emptyset). Since we have $2n$ players, we conclude $SW(\mathbf{X}(\mathbf{d}'), \mathbf{d}^*) \geq \frac{1}{c} SW(\mathbf{A}, \mathbf{d}^*)$, yielding the desired result. \square

Applying Lemma F.1 in place of Lemma 3.2, we can improve the statements of Theorems 3.7 and 4.1 so that the resulting prices of anarchy are improved from $c + 1$ to c , whenever algorithm \mathcal{A} is a c -approximation, but a $(c - 1)$ -approximation when agents are single-minded, and c is independent of n . This is the case, for example, in the standard greedy algorithm applied to cardinality-restricted combinatorial auctions.

Lemma F.2. *If \mathcal{A} is player-symmetric, object-symmetric, and a $c(n, m)$ -approximation, then for any declaration profile \mathbf{d} and allocation profile $\mathbf{A} = A_1, \dots, A_n$,*

$$\sum_{i \in [n]} d_i(X_i(\mathbf{d})) \geq \frac{1}{c(2n, 2m)} \sum_{i \in [n]} (\theta_i(A_i, \mathbf{d}_{-i}) + d_i(X_i(\mathbf{d}))).$$

Proof. Consider an auction with an additional copy of each player and each object; write i' for the copy of agent i , and M' for the additional objects. The feasibility constraints for the new objects and agents are identical to those for the original objects and agents. Then \mathcal{A} is a $c(2n, 2m)$ approximation algorithm for this new problem instance.

Choose any $\epsilon > 0$. We define \mathbf{d}' as in Lemma 3.2. We then define \mathbf{d}'' by setting $d_{i'}'' = d_i'$ and $d_{i''}''$ to be the single-minded declaration for set A_i at value $\theta_i(A_i, \mathbf{d}_{-i}) - \epsilon$. Finally, define \mathbf{d}''' by additionally adding a bid for the second copy of set $X_i(\mathbf{d})$ by agent i for value $d_i(X_i(\mathbf{d})) - \epsilon$. We then have $\mathbf{X}(\mathbf{d}''') = \mathbf{X}(\mathbf{d})$, but an alternative allocation gives A_i to each player i' , and the second copy of $X_i(\mathbf{d})$ to agent i . The result then follows since \mathcal{A} is a $c(2n, 2m)$ approximation. \square

Applying Lemma F.2 in place of Lemma 3.2, we can improve the statements of Theorems 3.7 and 4.1 so that the resulting price of anarchy is improved from $c + 1$ to c whenever the conditions of the Lemma apply and c is a constant.

We now give an example to show that these improved bounds are tight, even for pure Nash equilibria, for both $\mathcal{M}_1(\mathcal{A})$ and $\mathcal{M}_{crit}(\mathcal{A})$. That is, a pure Nash equilibrium can have approximation ratio as high as c for a c -approximate algorithm, where c is a constant, even if the algorithm is $(c - 1)$ -approximate when we assume that all bidders are single-minded.

Example F.3. Consider a combinatorial auction with the additional requirement that each bidder can be given at most 2 objects. The standard greedy algorithm that allocates in order of value is a 3 approximation. This algorithm and problem are player and object symmetric, and furthermore this algorithm is a 2 approximation when agents are single-minded.

Consider the following valuation profile. There are 3 bidders and 3 objects, say $\{a, b, c\}$. Choose arbitrarily small $\epsilon > 0$; the valuations of the players are as in the following table.

player	set	value
1	$\{a, b\}$	$1 + 3\epsilon$
1	$\{c\}$	1
2	$\{a\}$	1
2	$\{b, c\}$	$1 + \epsilon$
3	$\{b\}$	1

The optimal solution gives each player their desired singleton at a value of 1, for a total welfare of 3. However, one pure Nash equilibrium has each player bid truthfully, except having player 1 reduce his declared value for $\{a, b\}$ to the smallest value at which he will win it. This gives a total welfare of $1 + 3\epsilon$. So, for all of $\mathcal{M}_1(\mathcal{A})$, $\mathcal{M}_{crit}(\mathcal{A})$, and $\mathcal{M}_\mu(\mathcal{A})$, the price of anarchy is at least 3 in both pure and mixed strategies.

G Calculating Critical Prices

We note that some of our mechanisms require the calculation of critical prices. In many settings the calculation of critical prices is a simple task, but for some feasibility conditions the critical prices may not be obvious to extract. However, for a greedy allocation rule, the critical price for a given set can be determined to within an additive ϵ error in polynomial time by performing binary search.

Fix greedy allocation rule \mathcal{A} , agent i , set S , and declarations \mathbf{d}_{-i} . We wish to resolve the value of $\theta_i(S, \mathbf{d}_{-i})$ to within an additive error of ϵ . We perform the following binary search procedure on parameter $v \in \mathbf{R}$. We begin by setting $v = \epsilon$ and checking whether $X_i(d_i^v, \mathbf{d}_{-i}) = S$. If not, we double v and repeat. The first time $X_i(d_i^v, \mathbf{d}_{-i}) = S$, we stop doubling v and switch to applying binary search: If $X_i(d_i^v, \mathbf{d}_{-i}) = S$, decrease the value of v ; otherwise increase the value of v . This procedure resolves the value of v to within ϵ within $O(\log v_{max}/\epsilon)$ iterations. So the critical prices for all agents' allocated sets can be found in $O(n \log(v_{max}/\epsilon))$ invocations of algorithm \mathcal{A} .

H Proof of Theorem 5.1

We consider the proof of the first result; the others are similar. Our motivating intuition is that an agent will optimize his declarations for $\mathcal{M}_1(\mathcal{A}, \mathcal{A}')$ by optimizing for $\mathcal{M}_1(\mathcal{A})$ and $\mathcal{M}_1(\mathcal{A}')$ separately, as these mechanisms consider separate portions of the input. If this is true, then Theorem 3.4 immediately implies the desired result, as an equilibrium for $\mathcal{M}_1(\mathcal{A}, \mathcal{A}')$ must be a combination of an equilibrium for $\mathcal{M}_1(\mathcal{A})$ and an equilibrium for $\mathcal{M}_1(\mathcal{A}')$. The one concern in our intuition is that the desired strategies for $\mathcal{M}_1(\mathcal{A})$ and $\mathcal{M}_1(\mathcal{A}')$ may conflict. That is, a player may find it strategically advantageous to declare so that $d_i(M) < d_i(S)$ for some $S \subset M$. However, since we allow declarations to exhibit this non-monotonicity, declarations for $\mathcal{M}_1(\mathcal{A})$ and $\mathcal{M}_1(\mathcal{A}')$ are truly independent for all bidders and the result follows. \square

I Example omitted from Section 5

In this section we demonstrate that an allocation rule that applies a greedy algorithm and compares its outcome with allocating all objects to a single player, applying whichever outcome is best, may have unbounded price of anarchy under both the first-price and critical-price payment schemes.

Consider the combinatorial auction problem. Suppose \mathcal{A} is the non-adaptive greedy algorithm with priority rule $r(i, S, v) = v$ if $|S| \leq \sqrt{m}$, and $r(i, S, v) = 0$ otherwise. Let \mathcal{A}' be the non-adaptive greedy algorithm with priority rule $r(i, S, v) = v$ if $S = M$, and $r(i, S, v) = 0$ otherwise. Then \mathcal{A}' simply allocates the set of all objects to the player that declares the highest value for it. Let \mathcal{A}_{max} be the allocation rule that applies whichever of \mathcal{A} or \mathcal{A}' obtains the better result; that is, on input \mathbf{d} , \mathcal{A}_{max} returns $\mathbf{X}^{\mathcal{A}}(\mathbf{d})$ if $SW(\mathbf{X}^{\mathcal{A}}(\mathbf{d}), \mathbf{d}) > SW(\mathbf{X}^{\mathcal{A}'}(\mathbf{d}), \mathbf{d})$, otherwise returns $\mathbf{X}^{\mathcal{A}'}(\mathbf{d})$. It is known that \mathcal{A}_{max} is a $O(\sqrt{m})$ approximate algorithm [21].

Consider the following instance of the CA problem. We have $n = m \geq 2$, say with $M = \{a_1, \dots, a_m\}$. Choose $\epsilon > 0$ arbitrarily small. For each i , the private type of agent i , t_i , is the pointwise maximum of two single-minded valuation functions: one for set $\{a_i\}$ at value 1, and the other for set M at value $1 + \epsilon$. An optimal allocation profile for \mathbf{t} would assign $\{a_i\}$ to each agent i , for a total welfare of m .

Consider the following declaration profile. For each i , d_i is the single-minded valuation function for set M at value $1 + \epsilon$. On input \mathbf{d} , \mathcal{A}_{max} will assign M to some agent, for a total welfare of $1 + \epsilon$. Also, \mathbf{d} is a pure Nash equilibrium for $\mathcal{M}_1(\mathcal{A}_{max})$, $\mathcal{M}_{crit}(\mathcal{A}_{max})$, and $\mathcal{M}_\mu(\mathcal{A}_{max})$ for any μ : all agents receive a utility of 0, and there is no way for any single agent to obtain positive utility by deviating from \mathbf{d} . Taking $\epsilon \rightarrow 0$, we conclude that the price of anarchy for any of these mechanisms is $\Omega(m)$, which does not match the combinatorial $O(\sqrt{m})$ approximation ratio of \mathcal{A}_{max} .

J Applications

We now describe some applications of our results to particular combinatorial allocation problems, resulting in mechanisms whose prices of anarchy improve on the approximation ratios of the best known incentive compatible algorithms. Recall that we do not restrict agents to be single-minded, so known incentive compatible approximation algorithms for single-minded settings do not apply.

J.1 Combinatorial Auctions

The combinatorial auction problem is a blocking allocation problem. There is a greedy $\sqrt{2m}$ approximation algorithm for this problem [19]. By Theorem 3.4, the deterministic first-price mechanism for this algorithm has a $(\sqrt{2m} + O(\log m))$ Bayesian price of anarchy. Since the CA is a blocking allocation problem, this mechanism also has pure Nash equilibria, and its price of anarchy in pure strategies is $(\sqrt{2m} + 1)$.

An alternative allocation rule, which can be implemented with a polynomial number of demand queries, was proposed by Mu'alem and Nisan [21]. This allocation rule combines a standard greedy algorithm with an allocation of all objects to a single bidder. By Theorem 5.1, this algorithm can be implemented as a mechanism with $O(\sqrt{m})$ price of anarchy in mixed strategies. If we assume that agents will follow the weakly dominant set of strategies in which they do not overbid, $\mathcal{M}_{crit}(\mathcal{A})$ will always have a pure equilibrium and will have an $O(\sqrt{m})$ price of anarchy in pure strategies.

J.2 Cardinality-restricted Combinatorial Auctions

In the special case that players' desires are restricted to sets of size at most k , the standard greedy algorithm is k -approximate assuming single-minded agents. This translates to a $(k + 1)$ approximate algorithm for general agents, which we have shown can be implemented as a mechanism with a $(k + 1)$ price of anarchy assuming ex-post individually rational bidders (by Theorem 3.7 with Lemma F.1), or as a deterministic mechanism with a $k + O(\log k)$ price of anarchy with no such assumption (by Theorem 3.4). If $k \geq 2$ then this is a blocking allocation problem, and the first-price mechanism has a pure equilibrium and is a $(k + 1)$ approximation at any pure equilibrium by Theorem 4.1.

J.3 Multiple-Demand Unsplittable Flow Problem

Consider a variant of the unsplittable flow problem in which each agent has multiple terminal pairs, each with a different value, and wishes for one of them to be satisfied. An adaptive greedy algorithm obtains an $O(m^{\frac{1}{B-1}})$ approximation [7] for any $B > 1$, so Theorem 3.4 implies that the first-price mechanism for this algorithm yields a matching price of anarchy in mixed strategies.

J.4 Convex Bundle Auctions

In a convex bundle auction, M is the plane \mathbf{R}^2 , and allocations must be non-intersecting compact convex sets. We suppose that agents declare valuation functions by making bids for such sets. Given such a collection of bids, we define the aspect-ratio, R , to be the maximum diameter of a set divided by the minimum width of a set. A non-adaptive greedy allocation rule using a geometrically-motivated priority function yields an $O(R^{4/3})$ approximation [2]. Alternative greedy algorithms yield better approximation ratios for special cases, such as rectangles.

By Theorem 3.4, the deterministic first-price mechanism for this algorithm has a $O(R^{4/3})$ Bayesian price of anarchy. Since the CA is a blocking allocation problem, this mechanism has pure Nash equilibria, and its price of anarchy in pure strategies is also $O(R^{4/3})$.

J.5 Max-profit Unit Job Scheduling

In this problem, each bidder has a job of unit time to schedule on one of multiple machines. A bidder has various windows of time of the form (release time, deadline, machine) in which his job could be scheduled, with a potentially different profit resulting from each window. The profits and windows are private information to each bidder. The goal of the mechanism is to schedule the jobs to maximize the total profit. The standard greedy algorithm obtains a 3-approximation, and a 2-approximation when bidders are single-minded [20], and can thus be implemented as a mechanism that attains a price of anarchy of 3 assuming ex-post individually rational bidders. If we assume that agents will follow the weakly dominant set of strategies in which they do not overbid, $\mathcal{M}_{crit}(\mathcal{A})$ will always have a pure equilibrium and will have an price of anarchy of 3 in pure strategies.

Unlike the previous examples, an exact algorithm is known for the case of single-minded bidders [3], which uses dynamic programming and runs in time $O(n^7)$. We believe that this algorithm may be extended to handle k -minded bidders (i.e. where each valuation function is the maximum of at most k single-minded valuation functions), with a runtime of $O(n^7 k^7)$. Since this algorithm solves the problem optimally, it is incentive compatible. However, the greedy mechanism with price of anarchy 3 is still appealing since it is computationally efficient (as it runs in linear time) and easy for agents to understand and trust.