

Price of Anarchy for Greedy Auctions

B. Lucier*

A. Borodin†

Abstract

We study 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, and others. We focus on the problem of designing mechanisms that approximately optimize social welfare at every Bayesian-Nash equilibrium (BNE), which is the standard notion of equilibrium in settings of incomplete information. For a broad class of greedy approximation algorithms, we give a general black-box reduction to deterministic mechanisms with almost no loss to the approximation ratio at any BNE. We also consider the special case of Nash equilibria in full-information games, where we obtain tightened results. This solution concept is closely related to the well-studied price of anarchy. Furthermore, for a rich subclass of allocation problems, pure Nash equilibria are guaranteed to exist for our mechanisms. For many problems, the approximation factors we obtain at equilibrium improve upon the best known results for deterministic truthful mechanisms. In particular, we exhibit a simple deterministic mechanism for general combinatorial auctions that obtains an $O(\sqrt{m})$ approximation at every BNE.

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) 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 an-

swer to such issues is the development of approximation algorithms, but the VCG mechanism does not (in general) retain its truthfulness when applied to approximate solutions [26].

The incompatibility between approximations and standard mechanism design techniques has motivated the search for new, specially-tailored mechanisms for computationally intractable problems. This search has focused primarily on ex-post incentive compatible (IC) alternatives to the VCG mechanism. While this venture has been largely successful in settings where agent preferences are single-dimensional [1, 7, 20, 24], 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 [29].

In this extended abstract we study mechanisms that are not necessarily truthful, but rather yield good approximations at any equilibrium of agent behaviour. Such an equilibrium-based solution concept is standard in economic game theory. Performance at equilibrium has also been studied extensively in the algorithmic game theory literature as the *price of anarchy* of a given game: the ratio between the optimal outcome and the worst-case outcome at any equilibrium. Our approach is to apply these well-studied equilibrium notions directly to the field of algorithmic mechanism design.

A common equilibrium concept is that of (pure or mixed) Nash equilibrium (NE), whereby agents apply strategies (or distributions thereof) and no single agent has incentive to unilaterally deviate. The Nash equilibrium concept may be appropriate in some settings, such as repeated auctions (eg. for search engine advertising slots). However, in one-shot auctions, an assumption that agents bid at a Nash equilibrium is defensible only in full-information settings where all agent types are public knowledge. Modelling an auction as a full-information game seems unreasonable in most natural applications. We therefore consider an alternative that is suited to games of partial information: Bayesian Nash equilibrium (BNE). Bayesian equilibrium is the standard equilibrium concept in economics for games of incomplete information, proposed initially by Harsanyi [13]. Under this model, we suppose that the types of the

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

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

agents are drawn from known (not necessarily identical) distributions. After types are chosen, each agent applies a strategy that maximizes his expected 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. Note that any NE is also a BNE. We pose the question: can a given black-box approximation algorithm be converted into a mechanism that preserves its approximation ratio (up to first-order terms) at every BNE? We show that for a broad class of non-IC greedy algorithms, the answer is *yes*.

A few points require clarification. The concept of BNE requires that agents' types are drawn from commonly-known distributions. However, our mechanisms will not depend on the actual distributions themselves: we will present a single mechanism that works for *every* distribution. In economic terms, the mechanisms we consider are *detail-free*.

We note that the full-information concept of Nash equilibrium (NE) is a special case of BNE, so our mechanisms will also preserve approximation ratios at every (mixed or pure) Nash equilibria of a pure information game. This is precisely the concept of *price of anarchy* from the algorithmic game theory literature. Following Christodoulou et al [8], we can extend this notion to the *Bayesian price of anarchy*, which is the worst-case approximation attained at any BNE of the mechanism. Indeed, our motivating question can be rephrased as asking whether every c -approximation algorithm can be implemented as a mechanism with (Bayesian) price of anarchy $c(1 + o(1))$.

Dominant strategy truthfulness of an approximation mechanism is conceptually stronger as a solution concept than that of a mechanism that approximates the optimal social welfare at every equilibrium. However, as noted elsewhere [8], the NE and BNE solution concepts are not, strictly speaking, relaxations of dominant strategy truthfulness. There exist truthful mechanisms whose approximation ratios are not preserved at all Nash equilibria, such as the famous Vickrey auction.

While our mechanisms will be implementable in polynomial time (assuming an efficient implementation of the given approximation algorithm), we do not argue that a (Bayesian or non-Bayesian) equilibrium of the underlying game can necessarily be found in polynomial time. Analyzing how the participants in an auction would arrive at equilibrium (eg. in a repeated-auction setting) is left as an important open problem (which is partially addressed in a companion paper [22]).

1.1 Our Results 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. These algorithms are not generally incentive compatible [20]. Our first result is that if a greedy algorithm is paired with a first-price payment scheme (ie. 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 first-price mechanism that achieves a $(c + O(\log c))$ approximation at every mixed Bayesian Nash equilibrium.

We also show that the approximation ratio we obtain is tight (up to lower-order terms): there exist examples in which the first-price mechanism can have a $(c + \Omega(\log c))$ approximation at equilibrium.

Note that we prove our theorem for all *mixed* BNE, meaning that the approximation ratio holds even if agents can choose distributions over strategies. This is somewhat non-standard, as Bayesian equilibria are usually considered to include only pure strategies. Nevertheless, our result is more general than a consideration of only pure-strategy BNE, and also generalizes the (standard) notion of mixed (non-Bayesian) Nash equilibria.

Are there implementations of a c -approximate greedy algorithm for which the approximation ratio at Bayesian equilibria 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, 21], 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 in the presence of uncertainty. This possibility of overbidding can result in equilibria with poor social welfare. However, these bidding strategies are inherently risky: depending on the bids of other agents, an overbidding agent may end up with negative utility. If we can assume that agents do not participate in this risky behaviour, we can tighten

the approximation ratio we attain at equilibrium.

Theorem: Under the assumption that agents never declare more than their true values on any set, any greedy c -approximation algorithm for a combinatorial allocation problem can be implemented as a critical price mechanism that achieves a $(c + 1)$ approximation at every Bayesian Nash equilibrium.

We note that additional justifications for the “no-overbidding” assumption have appeared in the literature. Christodoulou et al [8] assume that no agent will bid in such a way that he might obtain negative utility (in their terms, the agents are “ex-post individually rational”), then define a mechanism so that an agent may pay his bid for any given set with some vanishingly small probability. Paes Leme and Tardos [21] justify a no-overbidding assumption by supposing that a random bid may appear in the input, again with vanishingly small probability. These approaches can be viewed as making assumptions about the risk-tolerance of bidders, then applying trembling-hand considerations in order to conclude that agents do not overbid. We could apply these trembling-hand techniques to our mechanisms as well, adding a factor of $(1 + \gamma)$ to our approximation ratios where γ can be made arbitrarily small. With this in mind, for the remainder of the paper we will simply state assumptions that agents do not overbid with little additional comment.

For certain algorithms, we can tighten our approximation ratios to 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 (ie. do not depend on agent or object labels). We discuss these improvements in Appendix A.

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, without any “no-overbidding” assumptions.

Theorem: Any greedy c -approximation algorithm for a combinatorial allocation problem can be implemented as a deterministic mechanism that achieves a $(c + 1)$ approximation at every pure Nash equilibrium.

The existence of an equilibrium in pure strategies is not guaranteed in general. We consider two important special cases for which we can guarantee the existence of pure equilibria (and retain our approximation ratio). First, we show that the critical-price mechanism for the *standard greedy algorithm*, which makes assignments in order of their value, always has a pure equilibrium. 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 (eg. for matroids, $k = 1$). Second, we present the class of *blocking allocation problems*, which essentially includes problems in which any agent can be allocated any given pair of objects. We show that a pure equilibrium is guaranteed to exist for a broad class of greedy allocation rules (those that are *non-adaptive* and *continuous*) for blocking allocation problems, paired with a first-price payment scheme.¹

Finally, we show how to extend our results for greedy algorithms to a combination of a greedy allocation rule with a rule that allocates all objects to a single bidder. Such combinations are a useful tool for constructing algorithms that can be implemented efficiently, such as for the general combinatorial auction problem [24].

1.2 Related Work The notion of Bayesian Nash equilibrium was introduced by Harsanyi [13]. For an overview of the development and impact of this theory we recommend a review by Myerson [25]. This and other equilibrium notions are common solution concepts for mechanism design in the economic literature; see Jackson [17] for a survey.

The inefficiency of equilibria is well-studied in the computational game theory literature, wherein the worst-case approximation ratio at equilibrium is referred to as the price of anarchy (introduced by Papadimitriou [28]). Inefficiency of equilibria is most commonly studied in settings in which agents choose their outcomes directly (eg. routing games [30]) 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 [27] and references therein.

The BNE solution concept has recently been 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. Paes Leme and Tardos [21] studied the performance of the generalized second price auction for advertising slots at equilibrium. Pure equilibria of first-price mechanisms have also been studied for path procurement auctions [16]. The problem of designing auctions that maximize revenue at Nash equilibrium has been extensively studied; notably in work on Internet advertising slot auctions [31, 11].

¹We will make the common assumption that valuation space is bounded and discretized by some arbitrarily small increment ϵ . The space of allowable types is then finite and infinitesimal utility improvements are precluded.

The most prominently studied allocation problem that falls into our framework is the combinatorial auction problem. Hastad's [14] result shows that it is NP-hard to approximate CAs to within $\Omega(m^{\frac{1}{2}-\epsilon})$ for any $\epsilon > 0$, even for succinctly representable valuation functions. The best known deterministic truthful mechanism for CAs with general valuations attains an approximation ratio of $O(\frac{m}{\sqrt{\log m}})$ [15]. A randomized $O(\sqrt{m})$ -approximate mechanism that is truthful in expectation was given by Lavi and Swamy [18]. 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 $O(Bm^{\frac{1}{B-2}})$ 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 ([7, 20, 24]) when bidders are restricted to being single-minded.

Our results make crucial use of the nature of greedy allocation algorithms. 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. The priority framework was extended to combinatorial auction problems by Borodin and Lucier [5]. Monotone greedy algorithms for combinatorial auctions were studied first by Lehmann et al [20], then subsequently by Mu'alem and Nisan [24] and 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* $v : 2^M \rightarrow \mathbb{R}$ 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. The *zero valuation* sets $v(S) = 0$ for

all $S \subseteq M$; we will represent this special valuation by \emptyset .

A *valuation profile* \mathbf{v} 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* \mathbf{X} 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* \mathcal{A} assigns to each valuation profile \mathbf{v} a feasible outcome $\mathcal{A}(\mathbf{v})$; we write $\mathcal{A}_i(\mathbf{v})$ for the allocation to agent i . We will tend to write \mathcal{A} for both an allocation rule and an algorithm that implements it.

Each agent $i \in [n]$ has a private valuation function t_i , his *type*, which defines the value he attributes to each allocation. 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 $SW(\mathcal{A}(\mathbf{t}), \mathbf{t}) \geq \frac{1}{c} SW_{opt}(\mathbf{t})$ 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 \mathcal{A} and a payment rule P . The mechanism proceeds by eliciting a valuation profile \mathbf{d} from each of the agents, called the *declaration profile*. 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 $\mathbf{d} \neq \mathbf{t}$. We will write $SW(\mathbf{d})$ for $SW(\mathcal{A}(\mathbf{d}), \mathbf{t})$ when the allocation rule and type profile are clear from context.

The utility of agent i in mechanism $\mathcal{M} = (\mathcal{A}, P)$, given declaration profile \mathbf{d} , is $u_i(\mathbf{d}) = t_i(\mathcal{A}_i(\mathbf{d})) - P_i(\mathbf{d})$. 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} .

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 $\mathbf{E}_{\mathbf{d} \sim \omega}[f(\mathbf{d})]$ 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 all $i \in [n]$ and distributions ω_i' ,

$$(2.1) \quad \mathbf{E}_{\mathbf{d} \sim \omega}[u_i(\mathbf{d})] \geq \mathbf{E}_{\mathbf{d} \sim (\omega_i', \omega_{-i})}[u_i(\mathbf{d})].$$

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

$$PoA_{mixed} = \sup_{\mathbf{t}, \omega} \frac{SW_{opt}(\mathbf{t})}{\mathbf{E}_{\mathbf{d} \sim \omega}[SW(\mathcal{M}(\mathbf{d}), \mathbf{t})]}$$

$$PoA_{pure} = \sup_{\mathbf{t}, \mathbf{d}} \frac{SW_{opt}(\mathbf{t})}{SW(\mathcal{M}(\mathbf{d}), \mathbf{t})}$$

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 Bayesian Types In a Bayesian setting, we suppose that the true types of the agents are not fixed, but are rather drawn from a known probability distribution \mathbf{D} over the set of valuation profiles. We assume that $\mathbf{D} = D_1 \times \dots \times D_n$ is the product of independent distributions, where $D_i(t_i)$ is the probability that agent i has type t_i . We write $SW_{opt}(\mathbf{D})$ for $\mathbf{E}_{\mathbf{t} \sim \mathbf{D}}[SW_{opt}(\mathbf{t})]$.

Given type t_i for agent i , let $d_i^{t_i}$ denote a declaration for agent i parameterized by t_i . We think of $d_i^{t_i}$ as the declaration that agent i will make if his true type is t_i . In a minor abuse of notation we will write $\mathbf{d}^{\mathbf{t}}$ for the declaration profile $(d_1^{t_1}, \dots, d_n^{t_n})$. A set of declaration profiles $\{\mathbf{d}^{\mathbf{t}}\}$ over all choices of \mathbf{t} forms a *Bayesian Nash Equilibrium* (BNE) if, for every i and every t_i in the support of D_i , agent i maximizes his expected utility by declaring $d_i^{t_i}$ whenever his type is t_i . That is, for all t_i and all d_i' ,

$$\mathbf{E}_{\mathbf{t}_{-i} \sim \mathbf{D}_{-i}}[u_i(d_i, \mathbf{d}_{-i}^{t_{-i}})] \geq \mathbf{E}_{\mathbf{t}_{-i} \sim \mathbf{D}_{-i}}[u_i(d_i', \mathbf{d}_{-i}^{t_{-i}})].$$

Assuming that the allocation rule is clear from context, we write $SW(\mathbf{D}, \{\mathbf{d}^{\mathbf{t}}\})$ for $\mathbf{E}_{\mathbf{t} \sim \mathbf{D}}[\sum_i t_i(\mathcal{A}_i(\mathbf{d}^{\mathbf{t}}))]$, the expected social welfare given types chosen from \mathbf{D} and strategies $\{\mathbf{d}^{\mathbf{t}}\}$.

We can generalize the notion of BNE to allow mixed types. Given 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 (randomized) 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 *mixed 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 $\omega_i^{t_i}$. That is, for each agent i , each possible type t_i , and every distribution ω_i' ,

$$\mathbf{E}_{\mathbf{t}_{-i} \sim \mathbf{D}_{-i}, \mathbf{d} \sim \omega^{\mathbf{t}}}[u_i(\mathbf{d})] \geq \mathbf{E}_{\mathbf{t}_{-i} \sim \mathbf{D}_{-i}, \mathbf{d} \sim (\omega_i', \omega_{-i}^{t_{-i}})}[u_i(\mathbf{d})].$$

We will write $SW(\mathbf{D}, \{\omega^{\mathbf{t}}\})$ to mean $\mathbf{E}_{\mathbf{t} \sim \mathbf{D}, \mathbf{d} \sim \omega^{\mathbf{t}}}[\sum_i t_i(\mathcal{A}_i(\mathbf{d}))]$, the expected social welfare given type distribution \mathbf{D} and strategy profiles $\omega^{\mathbf{t}}$.

The *mixed and pure Bayesian price of anarchy* of mechanism \mathcal{M} are defined as

$$BPoA_{mixed} = \sup_{\mathbf{D}, \{\omega^{\mathbf{t}}\}} \frac{SW_{opt}(\mathbf{D})}{SW_{\mathcal{M}}(\mathbf{D}, \{\omega^{\mathbf{t}}\})}$$

$$BPoA_{pure} = \sup_{\mathbf{D}, \{\mathbf{d}^{\mathbf{t}}\}} \frac{SW_{opt}(\mathbf{D})}{SW_{\mathcal{M}}(\mathbf{D}, \{\mathbf{d}^{\mathbf{t}}\})}$$

where the supremums are over all type distributions \mathbf{D} and mixed BNE $\{\omega^{\mathbf{t}}\}$ (respectively, pure BNE $\{\mathbf{d}^{\mathbf{t}}\}$) for \mathbf{D} .

2.3 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 [24], 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 N 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 N , some $i \notin N$, and allocation X_i , we say X_i is a *feasible allocation for i given N* if the partial allocation remains feasible when X_i is added to it.

A *monotone priority function* is a function $r : [n] \times 2^M \times \mathbb{R} \rightarrow \mathbb{R}$. We think of $r(i, S, v)$ as 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 set 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 S
4. Set $X_i = S$; remove player i from N
5. return X_1, \dots, X_n

We assume that ties in step 3 are broken in an arbitrary but fixed manner. 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 (eg. for submodular combinatorial auctions [19]). Such incremental allocation rules are not covered by our results; we leave open the BNE analysis of their implementations.

To build some intuition for our priority framework, we now mention a few examples of combinatorial allocation problems and greedy allocation rules. (See also

section 7.) The general combinatorial auction problem is defined by the feasibility constraint that no two allocations can intersect. Lehmann et al [20] show that the (non-adaptive) 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.

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 [23]. 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 . Each agent additionally specifies a fractional demand $d_i \in [0, 1]$ corresponding to a desired amount of flow to send along the given path. An allocation is feasible if the total allocated flow along each edge is no more than its capacity. 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.

2.4 Payment Methods Given allocation rule \mathcal{A} , 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 it, assuming that the

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 $O(Bm^{\frac{1}{B-2}})$ approximation algorithm due to Bartal et al [4].

other agents bid according to \mathbf{d}_{-i} . That is,

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

The *critical payment scheme* sets $P_i(\mathbf{d}) = \theta_i(\mathcal{A}_i(\mathbf{d}), \mathbf{d}_{-i})$, so each agent pays the critical price for the set she receives. We discuss implementation issues for this payment scheme in Section 6. By contrast, the first-price payment scheme sets $P_i(\mathbf{d}) = d_i(\mathcal{A}_i(\mathbf{d}))$, 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.

3 Mixed Bayesian 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.

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. An important property of a greedy algorithm is that 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.

DEFINITION 3.1. *An allocation rule \mathcal{A}' is loser-independent if, whenever \mathbf{d}_{-i} and \mathbf{d}'_{-i} satisfy $\mathcal{A}'(\emptyset, \mathbf{d}_{-i}) = \mathcal{A}'(\emptyset, \mathbf{d}'_{-i})$ and $d_j(\mathcal{A}'_j(\emptyset, \mathbf{d}_{-i})) = d_j(\mathcal{A}'_j(\emptyset, \mathbf{d}'_{-i}))$ for all $j \neq i$, then $\mathcal{A}'(d_i, \mathbf{d}_{-i}) = \mathcal{A}'(d_i, \mathbf{d}'_{-i})$ for all d_i .*

In other words, if \mathcal{A}' is a loser-independent algorithm, then agent i 's perception of the behaviour of \mathcal{A}' depends only on those agents who would win if agent i did not participate, and on their declared values for their winnings.

LEMMA 3.1. *\mathcal{A} is loser-independent.*

Proof. Choose $i \in [n]$ and $S \subseteq M$, and let \mathcal{A} be an adaptive greedy allocation rule. Choose \mathbf{d}_{-i} and \mathbf{d}'_{-i} such that $\mathcal{A}(\emptyset, \mathbf{d}_{-i}) = \mathcal{A}(\emptyset, \mathbf{d}'_{-i})$ and $d_j(\mathcal{A}_j(\emptyset, \mathbf{d}_{-i})) = d_j(\mathcal{A}_j(\emptyset, \mathbf{d}'_{-i}))$ for all $j \neq i$.

Recall the definition of an adaptive greedy algorithm, and consider the iterations of \mathcal{A} on input $(\emptyset, \mathbf{d}_{-i})$. 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 (assuming other agents declare according to \mathbf{d}_{-i}), 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 . Furthermore, for any d_i with $d_i(S) < v$, $\mathcal{A}_i(d_i, \mathbf{d}_{-i}) \neq S$. Thus $\theta_i(S, \mathbf{d}_{-i}) = v$. By the same argument we also have $\theta_i(S, \mathbf{d}'_{-i}) = v$, since \mathcal{A} allocates the same sets on inputs $(\emptyset, \mathbf{d}_{-i})$ and $(\emptyset, \mathbf{d}'_{-i})$ (by assumption) and in the same order (since $d_j(\mathcal{A}_j(\emptyset, \mathbf{d}_{-i})) = d_j(\mathcal{A}_j(\emptyset, \mathbf{d}'_{-i}))$ for all $j \neq i$, so the relative ranking order for all allocated sets is the same on every iteration). We conclude that $\theta_i(S, \mathbf{d}_{-i}) = \theta_i(S, \mathbf{d}'_{-i})$ for all $S \subseteq M$.

Choose some d_i and consider $\mathcal{A}_i(d_i, \mathbf{d}_{-i})$, the set allocated to agent i on declaration d_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 of \mathcal{A} . Next consider $\mathcal{A}_i(d_i, \mathbf{d}'_{-i})$. We know that $\theta_i(S, \mathbf{d}_{-i}) = \theta_i(S, \mathbf{d}'_{-i})$ for all S , and the behaviour of \mathcal{A} is identical on inputs (d_i, \mathbf{d}_{-i}) and (d_i, \mathbf{d}'_{-i}) up until the point at which an allocation to agent i is made. Thus the allocation to agent i will be the same on declarations (d_i, \mathbf{d}_{-i}) and (d_i, \mathbf{d}'_{-i}) , as required.

If \mathcal{A} is a c -approximate algorithm, then (on any input) the sum of the declared values for its output profile approximates the sum of the declared values for the optimal allocation. We now show that it also approximates 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 optimal allocation profile \mathbf{A} , $\sum_{i \in [n]} d_i(\mathcal{A}_i(\mathbf{d})) \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(\mathbf{A}_i, \mathbf{d}_{-i})$.*

Proof. Choose any $\epsilon > 0$. For all i , 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 . Since \mathcal{A} is loser-independent (by Lemma 3.1), the allocations of \mathcal{A} on inputs \mathbf{d} and \mathbf{d}^* are identical. Since \mathcal{A} is a c -approximation, we conclude that $SW(\mathcal{A}(\mathbf{d}), \mathbf{d}) = SW(\mathcal{A}(\mathbf{d}^*), \mathbf{d}^*) \geq \frac{1}{c} SW(\mathbf{A}, \mathbf{d}^*) \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(\mathbf{A}_i, \mathbf{d}_{-i}) - n\epsilon$. The result follows by taking the limit as $\epsilon \rightarrow 0$.

3.2 Greedy First-Price Mechanisms We now consider the performance of the first-price mechanism $\mathcal{M}_1(\mathcal{A})$ at equilibrium. We first note that the utility-maximizing declaration of an agent never involves over-bidding on a set that he may be allocated with positive probability, given any distribution over the declarations made by the other players.

LEMMA 3.3. *Suppose $\omega = (\omega_1, \dots, \omega_n)$ is a product distribution over declarations. If d_i maximizes the expected utility of agent i with respect to ω_{-i} , and $\Pr_{\mathbf{d}_{-i} \sim \omega_{-i}}[\mathcal{A}_i(d_i, \mathbf{d}_{-i}) = S] > 0$ for some $S \subseteq M$, then $d_i(S) \leq t_i(S)$.*

Proof. Suppose for contradiction that there exists some S such 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.

Fix any \mathbf{d}_{-i} in the support of ω_{-i} . We claim that $u_i(d_i', \mathbf{d}_{-i}) \geq u_i(d_i, \mathbf{d}_{-i})$. Let $S_i = \mathcal{A}_i(d_i', \mathbf{d}_{-i}), T_i = \mathcal{A}_i(d_i, \mathbf{d}_{-i})$. If $d_i(T_i) \leq t_i(T_i)$ then $u_i(d_i, \mathbf{d}_{-i}) \leq 0 \leq u_i(d_i', \mathbf{d}_{-i})$ as claimed. If, on the other hand, $d_i(T_i) > t_i(T_i)$, then $d_i'(T_i) = d_i(T_i)$ by definition. From the definitions of S_i and T_i , and of a priority algorithm, it must be that (on some iteration of \mathcal{A}) S_i has a higher priority than T_i under declaration d_i , but T_i has a higher priority than S_i under declaration d_i' . Since $d_i'(T_i) = d_i(T_i)$, $d_i'(S_i) \leq d_i(S_i)$, and r is monotone, this can occur only if $S_i = T_i$.³ Thus $\mathcal{A}_i(d_i', \mathbf{d}_{-i}) = \mathcal{A}_i(d_i, \mathbf{d}_{-i})$ and hence $u_i(d_i, \mathbf{d}_{-i}) = u_i(d_i', \mathbf{d}_{-i})$ as claimed.

We conclude that $u_i(d_i', \mathbf{d}_{-i}) \geq u_i(d_i, \mathbf{d}_{-i})$ for all \mathbf{d}_{-i} . Moreover, this inequality is strict when $\mathcal{A}_i(d_i, \mathbf{d}_{-i}) = S$, since $d_i(S) < t_i(S)$. Since this event occurs with positive probability, we conclude $\mathbf{E}_{\mathbf{d}_{-i} \sim \omega_{-i}}[u_i(d_i, \mathbf{d}_{-i})] < \mathbf{E}_{\mathbf{d}_{-i} \sim \omega_{-i}}[u_i(d_i', \mathbf{d}_{-i})]$, contradicting the maximality of d_i .

3.2.1 Warmup: Pure Nash Equilibria We are now ready to bound the price of anarchy of $\mathcal{M}_1(\mathcal{A})$. As a warmup, we will begin with a result for pure Nash equilibria, rather than the fully general BNE case. Note that the bound we obtain for pure equilibria is tighter than the bound we will obtain for the general case in Theorem 3.2.

THEOREM 3.1. *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 to the optimal social welfare.*

Proof. Fix type profile \mathbf{t} and suppose \mathbf{d} is a pure Nash equilibrium. Let A_1, \dots, A_n be an optimal allocation. Then

$$(3.2) \quad \sum_i t_i(\mathcal{A}_i(\mathbf{d})) \geq \sum_i d_i(\mathcal{A}_i(\mathbf{d})) \geq \frac{1}{c} \sum_i \theta_i(\mathbf{A}_i, \mathbf{d}_{-i})$$

where the first inequality follows from Lemma 3.3 and the second is Lemma 3.2.

³It is here that we make use of the fact that ties in rank are broken according to some arbitrary but fixed rule.

Choose arbitrarily small $\epsilon > 0$ and let d_i' be the single-minded declaration for set A_i at value $\theta_i(A_i, \mathbf{d}_{-i}) + \epsilon$. Then $\mathcal{A}_i(d_i', \mathbf{d}_{-i}) = A_i$ and hence $u_i(d_i', \mathbf{d}_{-i}) = t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i}) - \epsilon$. Since \mathbf{d} is a Nash equilibrium, it must be that

$$\begin{aligned} t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i}) - \epsilon &= u_i(d_i', \mathbf{d}_{-i}) \\ &\leq u_i(d_i, \mathbf{d}_{-i}) \\ &\leq t_i(\mathcal{A}_i(\mathbf{d})). \end{aligned}$$

Summing over all i and applying (3.2) we have

$$\sum_i t_i(\mathcal{A}_i(\mathbf{d})) \geq \frac{1}{c} \sum_i (t_i(A_i) - t_i(\mathcal{A}_i(\mathbf{d})) - \epsilon)$$

which, rearranging and taking $\epsilon \rightarrow 0$, implies

$$\begin{aligned} SW_{\mathcal{A}}(\mathbf{d}) &= \sum_i t_i(\mathcal{A}_i(\mathbf{d})) \\ &\geq \frac{1}{c+1} \sum_i t_i(A_i) \\ &= \frac{1}{c+1} SW_{opt} \end{aligned}$$

as required.

3.2.2 Bayesian Nash Equilibria We are now ready to prove our main result, which is a bound on the mixed Bayesian price of anarchy for mechanism $\mathcal{M}_1(\mathcal{A})$.

THEOREM 3.2. *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. Fix a distribution \mathbf{D} over type profiles and let $\{\omega^{\mathbf{t}}\}$ be a mixed Bayesian Nash equilibrium with respect to \mathbf{D} . Write τ for the distribution over declarations that results when types \mathbf{t} are chosen according to \mathbf{D} and declarations are then chosen according to $\omega^{\mathbf{t}}$. That is, for all i and d_i ,

$$\tau_i(d_i) = \sum_{\mathbf{t}_i} D_i(t_i) \omega_i^{t_i}(d_i).$$

For the remainder of the proof, unless otherwise specified, expectations over type profiles will be with respect to \mathbf{D} , and expectations over declaration profiles will be with respect to τ .

Choose some \mathbf{t} in the support of \mathbf{D} . Let $\mathbf{A}^{\mathbf{t}} = A_1^{\mathbf{t}}, \dots, A_n^{\mathbf{t}}$ denote an optimal allocation for \mathbf{t} . Following the proof of Theorem 3.1, we would like to bound the expected value of $\theta_i(A_i^{\mathbf{t}}, \mathbf{d}_{-i})$ with respect to $t_i(A_i^{\mathbf{t}})$ and $t_i(\mathcal{A}_i(\mathbf{d}))$ for each i . This will allow us to use Lemma 3.2 to obtain a relation between the expected welfare of \mathcal{A} and the expected optimal welfare. We encapsulate this bound in Claim 3.1.

CLAIM 3.1. *For all i , $S \subseteq M$ and $k > 1$, if d_i is in the support of $\omega_i^{t_i}$ then*

$$\begin{aligned} \mathbf{E}_{\mathbf{d}_{-i}}[\theta_i(S, \mathbf{d}_{-i})] &\geq \left(1 - \frac{1 + \log k}{k}\right) t_i(S) \\ &\quad - (1 + \log k) \mathbf{E}_{\mathbf{d}_{-i}}[t_i(\mathcal{A}_i(\mathbf{d}))]. \end{aligned}$$

Proof. For brevity, we will write $\bar{\theta}_i = \mathbf{E}_{\mathbf{d}_{-i}}[\theta_i(S, \mathbf{d}_{-i})]$ and $\bar{t}_i = \mathbf{E}_{\mathbf{d}_{-i}}[t_i(\mathcal{A}_i(\mathbf{d}))]$. That is, $\bar{\theta}_i$ is the expected critical price of set S for agent i , and \bar{t}_i is the expected welfare obtained by agent i when declaring d_i . Note that since d_i is in the support of $\omega_i^{t_i}$, which is assumed to be a strategy in equilibrium, d_i must maximize the expected utility of agent i , which is at most \bar{t}_i .

Suppose that $t_i(S) \leq \bar{t}_i$. In this case the result is trivially true, since

$$\begin{aligned} \bar{\theta}_i &\geq 0 \\ &\geq \left(1 - \frac{1 + \log k}{k}\right) (t_i(S) - \bar{t}_i) \\ &> \left(1 - \frac{1 + \log k}{k}\right) t_i(S) - (1 + \log k) \bar{t}_i. \end{aligned}$$

Next suppose that $\bar{t}_i = 0$. Then $\mathbf{E}_{\mathbf{d}_{-i}}[u_i(\mathbf{d})] = 0$. Suppose there exists some \mathbf{d}_{-i} in the support of τ_{-i} such that $\theta_i(S, \mathbf{d}_{-i}) < t_i(S)$. Then agent i would obtain positive expected utility by bidding single-mindedly for set S at value $\frac{1}{2}(\theta_i(S, \mathbf{d}_{-i}) + t_i(S))$. This contradicts the supposed optimality of d_i . It must therefore be that $\theta_i(S, \mathbf{d}_{-i}) \geq t_i(S)$ for all $\mathbf{d}_{-i} \sim \tau_{-i}$. We then have that

$$\bar{\theta}_i \geq t_i(S) \geq \left(1 - \frac{1 + \log k}{k}\right) t_i(S) - (1 + \log k) \bar{t}_i$$

as required.

Assume now that $t_i(S) > \bar{t}_i > 0$. Let $r = t_i(S)/\bar{t}_i$; then $r \in [1, \infty)$. For all $Z \geq 0$, let $p_Z = Pr_{\mathbf{d}_{-i}}[\theta_i(S, \mathbf{d}_{-i}) < t_i(S) - Z\bar{t}_i]$. Then we note that

$$(3.3) \quad \bar{\theta}_i \geq t_i(S) - \left(\int_0^\infty p_Z dZ\right) \bar{t}_i.$$

We claim that $p_Z \leq 1/Z$ for all $Z \in [1, r]$. Otherwise, if d_i' were the single-minded bid for set S at value $t_i(S) - Z\bar{t}_i$, then we would have

$$\mathbf{E}_{\mathbf{d}_{-i}}[u_i(d_i', \mathbf{d}_{-i})] = (Z\bar{t}_i)p_Z > \bar{t}_i \geq \mathbf{E}_{\mathbf{d}_{-i}}[u_i(d_i, \mathbf{d}_{-i})]$$

contradicting the optimality of d_i . We conclude that $p_Z \leq 1/Z$ for all $Z \in [1, r]$. Furthermore, $p_Z \leq 1$ for all $Z \leq 1$, trivially. Also, since we also know $\theta_i(S, \mathbf{d}_{-i}) \geq 0$ with probability 1, it must be that $p_Z = 0$ for all $Z > r$. Applying these bounds on p_Z to (3.3), we conclude

$$\bar{\theta}_i \geq t_i(S) - (1 + \log r) \bar{t}_i.$$

We now proceed by cases on the value of k .

Case 1: $r \leq k$. In this case

$$\bar{\theta}_i \geq t_i(S) - (1 + \log r)\bar{t}_i \geq t_i(S) - (1 + \log k)\bar{t}_i.$$

Case 2: $r \geq k$. In this case

$$\begin{aligned} \bar{\theta}_i &\geq t_i(S) - (1 + \log r)\bar{t}_i \\ &> \left(1 - \frac{1 + \log r}{r}\right) t_i(S) \\ &\geq \left(1 - \frac{1 + \log k}{k}\right) t_i(S). \end{aligned}$$

We conclude that, in all cases, either $\bar{\theta}_i \geq t_i(S) - (1 + \log k)\bar{t}_i$ or $\bar{\theta}_i \geq \left(1 - \frac{1 + \log k}{k}\right) t_i(S)$. This completes the proof of the claim.

Applying Claim 3.1 to set A_i^t , we conclude that for all $i \in [n]$, \mathbf{t} , and $k > 1$,

$$\begin{aligned} \mathbf{E}_{\mathbf{d}_{-i}}[\theta_i(A_i^t, \mathbf{d}_{-i})] &\geq \left(1 - \frac{1 + \log k}{k}\right) t_i(A_i^t) \\ &\quad - (1 + \log k)\mathbf{E}_{\mathbf{d}_{-i}}[t_i(\mathcal{A}_i(\mathbf{d}))]. \end{aligned}$$

Summing over i and taking expectation over all choices of \mathbf{t} and d_i , we have

$$(3.4) \quad \begin{aligned} \mathbf{E}_{\mathbf{t}, \mathbf{d}} \left[\sum_i \theta_i(A_i^t, \mathbf{d}_{-i}) \right] \\ \geq \left(1 - \frac{1 + \log k}{k}\right) \mathbf{E}_{\mathbf{t}} \left[\sum_i t_i(A_i^t) \right] \\ - (1 + \log k) \sum_i \mathbf{E}_{\mathbf{t}, \mathbf{d}_{-i}, d_i \sim \omega_i^{t_i}} [t_i(\mathcal{A}_i(\mathbf{d}))]. \end{aligned}$$

Note that $\mathbf{E}_{\mathbf{t}}[\sum_i t_i(A_i^t)]$ is precisely $SW_{opt}(\mathbf{D})$. Additionally,

$$\begin{aligned} \sum_i \mathbf{E}_{\mathbf{t}, \mathbf{d}_{-i}, d_i \sim \omega_i^{t_i}} [t_i(\mathcal{A}_i(\mathbf{d}))] &= \sum_i \mathbf{E}_{\mathbf{t}, \mathbf{d} \sim \omega^t} [t_i(\mathcal{A}_i(\mathbf{d}))] \\ &= SW(\mathbf{D}, \{\omega^t\}). \end{aligned}$$

Finally,

$$\begin{aligned} \mathbf{E}_{\mathbf{t}, \mathbf{d}} \left[\sum_i \theta_i(A_i^t, \mathbf{d}_{-i}) \right] \\ \leq c\mathbf{E}_{\mathbf{t}, \mathbf{d}} \left[\sum_i d_i(\mathcal{A}_i(\mathbf{d})) \right] \quad (\text{Lemma 3.2}) \\ = c\mathbf{E}_{\mathbf{t}, \mathbf{d} \sim \omega^t} \left[\sum_i d_i(\mathcal{A}_i(\mathbf{d})) \right] \\ \leq c\mathbf{E}_{\mathbf{t}, \mathbf{d} \sim \omega^t} \left[\sum_i t_i(\mathcal{A}_i(\mathbf{d})) \right] \quad (\text{Lemma 3.3}) \\ = cSW(\mathbf{D}, \{\omega^t\}). \end{aligned}$$

We conclude from (3.4) that

$$\begin{aligned} cSW(\mathbf{D}, \{\omega^t\}) &\geq \left(1 - \frac{1 + \log k}{k}\right) SW_{opt}(\mathbf{D}) \\ &\quad - (1 + \log k)SW(\mathbf{D}, \{\omega^t\}). \end{aligned}$$

Setting $k = c$ and rearranging yields

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

as required.

We next show by way of an example that the analysis in Theorem 3.2 is tight.

PROPOSITION 3.1. *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} , and the mixed price of anarchy for $\mathcal{M}_1(\mathcal{A})$ is $c + \Omega(\log c)$.*

Proof. Our problem will be 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 $\{a_{11}, a_{21}, \dots, a_{i1}\}$ 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 $\frac{1}{2} \cdot \frac{2}{3} \cdots \frac{X-1}{X} = \frac{1}{X} \leq \frac{1}{Z}$. 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 $c + \Omega(\log c)$.

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 3.3, which plays a crucial part in the proof of Theorem 3.2, fails to hold for the critical payment rule: there is no guarantee that rational 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 following example shows.

Example. Consider a combinatorial auction with 3 objects, $\{a, b, c\}$, and 3 bidders, under the feasibility 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 analysis 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. 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.

We get around the issue of overbidding by directly assuming that agents do not overbid. Such an assumption is most reasonable in settings where agents can be presumed to be averse to risking negative utility. Given that agents will not overbid, a simple modification of Theorem 3.2 yields a sharpened result.

THEOREM 3.3. *Suppose \mathcal{A} is a c -approximate greedy allocation rule. Suppose also that bidders do not overbid. Then the expected welfare at any mixed Bayesian Nash equilibrium of $\mathcal{M}_{crit}(\mathcal{A})$ is a $(c + 1)$ approximation to the optimal welfare.*

Proof. For brevity and clarity we will limit the following proof to mixed Nash equilibria (ie. non-Bayesian types). The extension to Bayesian types is straightforward but notationally cumbersome.

Fix a true type profile \mathbf{t} , and let ω be a mixed Nash equilibrium with respect to \mathbf{t} . Let A_1, \dots, A_n denote an optimal allocation for \mathbf{t} . Lemmas 3.3 and 3.2 imply that, for all $\mathbf{d} \sim \omega$,

$$\sum_i t_i(\mathcal{A}_i(\mathbf{d})) \geq \frac{1}{c} \sum_i \theta_i(A_i, \mathbf{d}_{-i}).$$

Taking expectation over declaration profiles with respect to distribution ω , we have that

$$(3.5) \quad \mathbf{E}_{\mathbf{d} \sim \omega} \left[\sum_i t_i(\mathcal{A}_i(\mathbf{d})) \right] \geq \frac{1}{c} \mathbf{E}_{\mathbf{d} \sim \omega} \left[\sum_i \theta_i(A_i, \mathbf{d}_{-i}) \right].$$

Let d_i' be the single-minded declaration for set A_i at value $t_i(A_i)$. We have

$$\begin{aligned} & \mathbf{E}_{\mathbf{d}_{-i} \sim \omega_{-i}} [u_i(d_i', \mathbf{d}_{-i})] \\ & \geq \sum_{\mathbf{d}_{-i}} \omega_{-i}(\mathbf{d}_{-i}) (t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i})) \\ & = t_i(A_i) - \mathbf{E}_{\mathbf{d}_{-i} \sim \omega_{-i}} [\theta_i(A_i, \mathbf{d}_{-i})]. \end{aligned}$$

Since ω_i maximizes the expected utility of agent i , we must have that

$$\begin{aligned} t_i(A_i) - \mathbf{E}_{\mathbf{d}_{-i} \sim \omega_{-i}} [\theta_i(A_i, \mathbf{d}_{-i})] & \leq \mathbf{E}_{\mathbf{d} \sim \omega} [u_i(\mathbf{d})] \\ & \leq \mathbf{E}_{\mathbf{d} \sim \omega} [t_i(\mathcal{A}_i(\mathbf{d}))]. \end{aligned}$$

Summing over all i and applying (3.5), we conclude

$$\begin{aligned} & \mathbf{E}_{\mathbf{d} \sim \omega} \left[\sum_i t_i(\mathcal{A}_i(\mathbf{d})) \right] \\ & \geq \frac{1}{c} \left(\sum_i t_i(A_i) - \mathbf{E}_{\mathbf{d} \sim \omega} \left[\sum_i t_i(\mathcal{A}_i(\mathbf{d})) \right] \right) \end{aligned}$$

which implies

$$\mathbf{E}_{\mathbf{d} \sim \omega} [SW_{\mathcal{A}}(\mathbf{d})] \geq \frac{1}{c+1} SW_{opt}(\mathbf{t})$$

as required.

The bound in Theorem 3.3 can be sharpened further to c in certain important special cases. This follows directly from a corresponding improvement to Lemma 3.2, discussed in Appendix A.

4 Existence of Pure Nash Equilibria

In this section we discuss pure equilibria in the full-information game for a mechanism with a greedy allocation rule. Recall from Theorem 3.1 that the approximation factor from Theorem 3.2 is improved from $(c + O(\log c))$ to $(c + 1)$ when we restrict our attention to equilibria in pure strategies. However, the power of Theorem 3.1 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. This is true even under the assumption that private valuations and payments are discretized, so that all values and payments are multiples of some arbitrarily small $\epsilon > 0$. A simple example is given below.

Example. Consider an instance of the combinatorial auction problem with two objects, $M = \{a, b\}$, and three agents. Each agent can be assigned at most one object, and moreover agent 2 cannot be allocated object b , and agent 3 cannot be allocated object a . Let

\mathcal{A} be the standard greedy algorithm. 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 Nash equilibrium exists for this example, even assuming that all private types and payments are multiples of some $\epsilon > 0$. Assume for contradiction that there is a Nash equilibrium \mathbf{d} for type profile \mathbf{t} and mechanism $\mathcal{M}_1(\mathcal{A})$.

We know that agent 1 does not win item b with a payment greater than 2, as this would cause him negative utility (and thus he would not be in equilibrium). Thus it must be that $\mathcal{A}_3(\mathbf{d}) = \{b\}$, 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 $\mathcal{A}_1(\mathbf{d}) = \{a\}$, 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, since agent 3 maximizes his utility up to an additive ϵ , 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.

4.1 Pure Equilibria for First-Price Mechanisms

Motivated by the example given above, we will consider special cases in which we can guarantee the existence of pure equilibria. We first consider a restriction on the class of allocation problems. Such a restriction amounts to 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 4.1. (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 declares a large enough value for them) is a blocking allocation problem. We will additionally assume that valuation space is discretized by some arbitrarily small increment $\epsilon > 0$, so that agents do not differentiate between outcomes whose

utilities are within an additive difference of ϵ . Note that the price of anarchy bound from Theorem 3.1 holds in such a discretized setting, within an additive error on total social welfare 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 .⁴ Many natural greedy algorithms are continuous; see Section 7 for examples. We now show that $\mathcal{M}_1(\mathcal{A})$ has a pure Nash equilibrium for any blocking allocation problem, when \mathcal{A} is non-adaptive and continuous.

THEOREM 4.1. *Suppose \mathcal{A} is a c -approximate non-adaptive continuous greedy allocation rule for a blocking allocation problem. Then $\mathcal{M}_1(\mathcal{A})$ has a pure Nash equilibrium.*

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

Define $r_i := r(i, \mathcal{A}_i(\mathbf{t}), t_i(\mathcal{A}_i(\mathbf{t})))$ for each agent i . Let ℓ be the agent with minimal r_ℓ (breaking ties arbitrarily), and let d_ℓ be the single-minded declaration for set $\mathcal{A}_\ell(\mathbf{t})$ at value ϵ . For every other agent $i \neq \ell$, let d_i be the single-minded declaration for set $\mathcal{A}_i(\mathbf{t})$ at value $t_i(\mathcal{A}_i(\mathbf{t}))$.

For every pair of agents i, j such that $r_i > r_j$, and every set S such that $\mathcal{A}_i(\mathbf{t}) \not\subseteq S$ and $t_i(S) \geq \theta_i(S, \mathbf{d}_{-i})$, find a set R_{ijS} such that R_{ijS} is a feasible allocation to agent j given that each agent k with $r_k > r_i$ is allocated $\mathcal{A}_k(\mathbf{t})$, and additionally no outcome allocates R_{ijS} to bidder j and either S or $\mathcal{A}_i(\mathbf{t})$ to bidder i . Such an R_{ijS} must exist from the definition of a blocking allocation problem, and moreover $R_{ijS} \not\subseteq \mathcal{A}_j(\mathbf{t})$ (since there *does* exist a feasible allocation that allocates $\mathcal{A}_j(\mathbf{t})$ to agent j and $\mathcal{A}_i(\mathbf{t})$ to agent i , namely $\mathcal{A}(\mathbf{t})!$). Additionally choose value v_{ijS} so that $r(i, S, t_i(S) - \epsilon) \leq r(j, R_{ijS}, v_{ijS}) < r(i, S, t_i(S))$. Such a value of v_{ijS} must exist, since \mathcal{A} is a continuous algorithm with bounded approximation ratio. Let d_{ijS} be the single-minded declaration for set R_{ijS} with value v_{ijS} . Note then that $r(j, R_{ijS}, v_{ijS}) < r_i$ (since $r(i, S, t_i(S)) \leq$

⁴Note that we do not require in this definition that v be a multiple of ϵ .

$r(i, \mathcal{A}_i(\mathbf{t}), t_i(\mathcal{A}_i(\mathbf{t})))$, so $v_{ijS} < \theta_j(R_{ijS}, \mathbf{d}_{-j})$. Define $d_{j'}$ to be the pointwise maximum of d_j and d_{ijS} for all i and S as described above.

We have finished the definition of \mathbf{d}' . Note that since all of the single-minded bids making up declarations d_{ijS} fall below their critical prices, $\mathcal{A}(\mathbf{d}') = \mathcal{A}(\mathbf{d}) = \mathcal{A}(\mathbf{t})$ by the loser-independence of \mathcal{A} . We claim that no agent can improve his utility by ϵ with a deviation from \mathbf{d}' . First consider agent ℓ . Under declaration profile \mathbf{d}' , agent ℓ obtains set $\mathcal{A}_\ell(\mathbf{t})$ for a price of ϵ . Since r_ℓ is minimal, $\theta_i(S, \mathbf{d}'_{-i}) \geq t_i(S)$ for all S such that $\mathcal{A}_\ell(\mathbf{t}) \not\subseteq S$. Next consider any agent $i \neq \ell$. For any set S such that $\mathcal{A}_i(\mathbf{t}) \not\subseteq S$, $\theta_i(S, \mathbf{d}'_{-i}) \geq t_i(S) - \epsilon$, due to the presence of a bid for set R_{ijS} by an agent j . Thus no change in declaration from \mathbf{d}' can improve the utility of agent i by more than ϵ , and thus cannot improve utility of agent i under our discretization assumptions. We conclude that \mathbf{d}' forms a pure Nash equilibrium, as required.

4.2 Pure Equilibria for Critical-Price Mechanisms Theorem 4.1 is restricted to the scope of blocking allocation problems. In this section we consider an alternative pure equilibrium construction for arbitrary (blocking and 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. We would like to guarantee the existence of such equilibria, but also retain the price of anarchy result from Theorem 3.1. To do this, we will require the assumption from Section 3.3 that agents do not overbid. We note that this assumption is somewhat better motivated in a pure equilibrium setting: as we show below, if an agent knows the bids to be made by the other bidders, that agent cannot gain from overbidding. Compare this to the example in Section 3, where it was shown that overbidding may be beneficial in settings of uncertainty.

CLAIM 4.1. *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 $S \subseteq M$.*

Proof. Choose any utility-maximizing d_i' , say with $\mathcal{A}_i(d_i', \mathbf{d}_{-i}) = T$. Let d_i be the single-minded declaration for set T with value $t_i(T)$; then $d_i(S) \leq t_i(S)$ for all S by the monotonicity of t_i . Since d_i' is utility-maximizing, we know that

$$0 \leq u_i(d_i', \mathbf{d}_{-i}) = t_i(A_i) - \theta_i(A_i, \mathbf{d}_{-i}).$$

So $t_i(A_i) \geq \theta_i(A_i, \mathbf{d}_{-i})$ and hence $\mathcal{A}_i(d_i, \mathbf{d}_{-i}) = T$. We therefore have $u_i(d_i', \mathbf{d}_{-i}) = u_i(d_i, \mathbf{d}_{-i})$, since $\mathcal{M}_{crit}(\mathcal{A})$ uses the critical pricing scheme. Thus, since d_i' is utility-maximizing, d_i is as well.

We now show that if \mathcal{A} is the standard greedy algorithm, then there is a pure Nash equilibrium for $\mathcal{M}_{crit}(\mathcal{A})$. Moreover, if agents do not overbid, then the mechanism achieves a $(c+1)$ approximation at any pure equilibrium.

THEOREM 4.2. *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. Moreover, if agents do not overbid, then $\mathcal{M}_{crit}(\mathcal{A})$ obtains a $(c+1)$ approximation to the optimal social welfare at every pure Nash equilibrium.*

Proof. Define \mathbf{d} by having d_i be the single-minded bid for set $\mathcal{A}_i(\mathbf{t})$ at value $t_i(\mathcal{A}_i(\mathbf{t}))$. Then the critical price for each agent's winning set is 0, so $u_i(\mathbf{d}) = t_i(\mathcal{A}_i(\mathbf{t}))$ for each i . Consider any agent i and any set T . Since \mathcal{A} ranks by value, either $t_i(T) \leq t_i(\mathcal{A}_i(\mathbf{t}))$ or else $\theta_i(T, \mathbf{t}_{-i}) > t_i(T)$. In either case, for any d_i such that $\mathcal{A}_i(d_i, \mathbf{t}_{-i}) = T$, it must be that $u_i(d_i, \mathbf{t}_{-i}) \leq u_i(\mathbf{d})$. We conclude that \mathbf{t} forms an equilibrium.

The price of anarchy result follows precisely the proof of Theorem 3.1, applying the assumption that no agent will overbid in place of Lemma 3.3.

5 Combining Mechanisms

A standard 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 [4, 7, 24]. When bidders are single-minded, such a combination rule will be incentive-compatible [24]. 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. Consider the following example.

Example. 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 $\mathcal{A}(\mathbf{d})$ if $SW(\mathcal{A}(\mathbf{d}), \mathbf{d}) >$

$SW(\mathcal{A}'(\mathbf{d}), \mathbf{d})$, otherwise returns $\mathcal{A}'(\mathbf{d})$. It is known that \mathcal{A}_{max} is a $O(\sqrt{m})$ approximate algorithm [24].

Our instance of the CA problem is the following. 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 .

We construct a declaration profile as follows. 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} .

In light of the above example, we 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 (eg. combinatorial auctions, see Section 7.1) the resulting randomized allocation rule obtains (in expectation) 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 of generality for many algorithms of interest (eg. combinatorial auctions, again see Section 7.1). 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 $\mathcal{M}_{crit}(\mathcal{A}, \mathcal{A}')$ 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 monotone 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 BNE.*

Proof. Since the portions of agent declarations relevant to $\mathcal{M}_1(\mathcal{A})$ and $\mathcal{M}_1(\mathcal{A}')$ are independent, an agent will optimize his declaration for $\mathcal{M}_1(\mathcal{A}, \mathcal{A}')$ by optimizing for $\mathcal{M}_1(\mathcal{A})$ and $\mathcal{M}_1(\mathcal{A}')$ separately. Theorem 3.2 then 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 $\mathcal{M}_1(\mathcal{A}')$.

6 Calculating Critical Prices

We note that some of our mechanisms require the calculation of critical prices. For many allocation algorithms, the calculation of critical prices is a simple task, which can be performed in parallel with the computation of an allocation profile. We leave the development of such pricing methods to the creators of the allocation algorithms to which our reduction may be applied. However, even if a specially-tailored algorithm for computing exact critical prices is not available, we note that critical prices for a given black-box greedy algorithm can be determined to within an additive ϵ error in polynomial time via simple binary search. Thus, assuming that valuation space is discretized, critical prices can be determined efficiently.

We now describe the procedure for determining critical prices in more detail. Fix greedy allocation rule \mathcal{A} , agent i , and declarations \mathbf{d} . Write v_{max} for $\max_{i,S} d_i(S)$. Suppose that $\mathcal{A}_i(d_i, \mathbf{d}_{-i}) = S$. Assume further that the space of possible values for sets is discretized by increments of some $\epsilon > 0$. 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 $v \in \mathbb{R}$. Write d_i^v for the single-minded declaration for set S at value v . Begin by setting $v = \epsilon$ and checking whether $\mathcal{A}_i(d_i^v, \mathbf{d}_{-i}) = S$. If not, double v and repeat. The first time $\mathcal{A}_i(d_i^v, \mathbf{d}_{-i}) = S$, stop doubling v and switch to applying binary search: If $\mathcal{A}_i(d_i^v, \mathbf{d}_{-i}) = S$, decrease the value of v ; otherwise increase the value of v . This procedure resolves the value of v to within ϵ in $O(\log v_{max}/\epsilon)$ iterations. Thus, for any given input to mechanism $\mathcal{M}_{crit}(\mathcal{A})$, the critical prices for all agents' allocated sets can be found in $O(n \log(v_{max}/\epsilon))$ invocations of algorithm \mathcal{A} . Note that it is not necessary for v_{max} to be known by the mechanism, since the mechanism need only calculate critical prices for sets that are actually allocated on input \mathbf{d} , and thus the maximum value declared in \mathbf{d} serves as an upper bound for critical prices.

7 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 compat-

ible 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.

7.1 Combinatorial Auctions The combinatorial auction problem is a blocking allocation problem. There is a greedy non-adaptive $\sqrt{2m}$ approximation algorithm for this problem [20]. By Theorem 3.2, 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 [24]. 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 $O(\sqrt{m})$ Bayesian price of anarchy.

7.2 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 can be implemented as a mechanism with a $(k + 1)$ price of anarchy assuming that agents do not overbid (by Theorem 3.3 with Lemma A.1), or as a mechanism with a $k + O(\log k)$ Bayesian price of anarchy with no such assumption (by Theorem 3.2). If $k \geq 2$ then this is a blocking allocation problem, and the first-price mechanism has a pure equilibrium and obtains a $(k + 1)$ approximation at any pure equilibrium by Theorem 3.1.

7.3 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 $O(Bm^{\frac{1}{B-1}})$ approximation [7] for any $B > 1$, so Theorem 3.2 implies that the first-price mechanism for this algorithm yields a matching price of anarchy in mixed strategies.

7.4 Convex Bundle Auctions In a convex bundle auction, M is the plane \mathbb{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, the aspect-ratio, R , is defined 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 3.2, the deterministic first-price mechanism for this algorithm has a $O(R^{4/3})$ Bayesian price of anarchy. Since this auction is a blocking allocation problem, this mechanism has pure Nash equilibria, and its price of anarchy in pure strategies is also $O(R^{4/3})$.

7.5 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 [23], and can thus be implemented as a mechanism that attains a price of anarchy of 3 assuming that bidders do not overbid. 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 (ie. 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.

8 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 Bayesian Nash equilibrium or any mixed Nash equilibrium. This has a number of applications, such as a combinatorial auction with $O(\sqrt{m})$ price of anarchy. 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 3.3; 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 4.1 and 4.2, finding other classes of problems for which pure equilibria are guaranteed to exist.

While we have studied performance at equilibrium, we have not addressed the issue of how agents should arrive at such an equilibrium. One desirable property along these lines would be a mechanism with only a single equilibrium, which can be easily computed. Additionally, one might model bidder behaviour via response dynamics, and ask whether agents are expected to converge to an equilibrium in a repeated-auction setting. This question is partially addressed in a related paper [22] that considers the design of mechanisms that account for bidder behaviour in repeated-auction settings.

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 a black-box 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. We are especially thankful to the anonymous referees whose input greatly improved the quality of this paper.

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. C. Harsanyi. Games with incomplete information played by ‘bayesian’ players, parts i ii and iii. *Management Science*, 14, 1967-68.
- [14] J. Hastad. Some optimal inapproximability results. In *Proc. 29th ACM Symp. on Theory of Computing*, 1997.
- [15] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.
- [16] N. Immorlica, D. Karger, E. Nikolova, and R. Sami. First-price path auctions. In *Proc. 7th ACM Conf. on Electronic Commerce*, 2005.
- [17] M. Jackson. A crash course in implementation theory. In *Social Choice and Welfare*, 2001.
- [18] 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.
- [19] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proc. 3rd ACM Conf. on Electronic Commerce*, 2001.
- [20] 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.
- [21] R. Paes Leme and E. Tardos. Sponsored search equilibria for conservative bidders. In *Fifth Workshop on Ad Auctions*, 2009.
- [22] B. Lucier. Beyond equilibria: Mechanisms for repeated combinatorial auctions. Working Paper, 2009.
- [23] J. Mestre. Greedy in approximation algorithms. In *Proc. 14th European Symp. on Algorithms*, 2006.
- [24] A. Mu’alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64:612–631, 2008.
- [25] R. Myerson. Harsanyi’s games with incomplete information. In *Management Science*, 2004.
- [26] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. 31st ACM Symp. on Theory of Computing*, pages 129–140. ACM Press, 1999.
- [27] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [28] C. Papadimitriou. Algorithms, games and the internet. In *Proc. 33rd ACM Symp. on Theory of Computing*, pages 749–752. ACM Press, 2001.
- [29] C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, 2008.
- [30] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 2002.
- [31] Hal Varian. Position auctions. Working Paper, UC Berkeley, 2006.

A 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 3.3 and 3.1. We say that a combinatorial allocation problem is *player symmetric* if the feasibility constraints do not depend on the labels of the players, and *object symmetric* if they do not depend on the labels of the objects. We say that a greedy algorithm is *player symmetric* if its ranking function r does not depend on its first parameter, and we say that it is *object symmetric* if its ranking function r does not distinguish between sets of the same cardinality in its second parameter.

LEMMA A.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 $\mathbf{A} = A_1, \dots, A_n$, $\sum_{i \in [n]} d_i(\mathcal{A}_i(\mathbf{d})) \geq \frac{1}{c(2n)} \sum_{i \in [n]} \theta_i(A_i, \mathbf{d}_{-i})$*

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 $\mathcal{A}(\mathbf{d}') = \mathcal{A}(\mathbf{d}'')$ (meaning that each additional player is allocated \emptyset). Since we have $2n$ players, we conclude $SW(\mathcal{A}(\mathbf{d}'), \mathbf{d}^*) \geq \frac{1}{c} SW(\mathbf{A}, \mathbf{d}^*)$, yielding the desired result.

Applying Lemma A.1 in place of Lemma 3.2, we can improve the statements of Theorems 3.3 and 3.1 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 A.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 $\mathbf{A} = A_1, \dots, A_n$,*

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

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 $\mathcal{A}_i(\mathbf{d})$ by agent i for value $d_i(\mathcal{A}_i(\mathbf{d})) - \epsilon$. We then have $\mathcal{A}(\mathbf{d}''') = \mathcal{A}(\mathbf{d})$, but an alternative allocation gives A_i to each player i' , and the second copy of $\mathcal{A}_i(\mathbf{d})$ to agent i . The result then follows since \mathcal{A} is a $c(2n, 2m)$ approximation.

Applying Lemma A.2 in place of Lemma 3.2, we can improve the statements of Theorems 3.3 and 3.1 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. 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.