

Equilibria of Greedy Combinatorial Auctions

BRENDAN LUCIER, Microsoft Research
ALLAN BORODIN, University of Toronto

We consider auctions in which greedy algorithms, paired with first-price or critical-price payment rules, are used to resolve multi-parameter combinatorial allocation problems. We study the price of anarchy for social welfare in such auctions. We show for a variety of equilibrium concepts, including Bayes-Nash equilibrium, low-regret bidding sequences, and asynchronous best-response dynamics, the resulting price of anarchy bound is close to the approximation factor of the underlying greedy algorithm.

1. INTRODUCTION

The field of algorithmic mechanism design studies systems that depend upon interaction with participants whose behaviour is motivated by their own goals, rather than those of a designer. Relevant solutions must therefore merge the computational considerations of computer science with the game-theoretic insights of economics. The focus of this paper is the multi-parameter domain of combinatorial allocation problems when the goal is to assign m objects to n agents in order to maximize the social welfare, subject to arbitrary downward-closed feasibility constraints. This class includes all combinatorial auction problems that allow single-minded declarations including multi-unit combinatorial auctions, unsplittable flow problems, and many others.

For the goal of optimizing social welfare, 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-maximization problem be solved exactly. For all but the simplest settings, this optimality requirement is undesirable: exact maximization may be computationally intractable, it may require an unrealistic amount of communication from the buyers, and the resulting winner determination rules may be difficult to explain to a typical participant. One way to bypass these complexity issues is to design new, specially-tailored mechanisms for specific assignment problems. Indeed, there has been significant progress in designing dominant strategy incentive compatible (DSIC) alternatives to the VCG mechanism. While this venture has been largely successful in settings where agent preferences are single-dimensional [Archer and Tardos 2001; Briest et al. 2005; Lehmann et al. 1999; Mu'alem and Nisan 2008], general settings have proven more difficult. It has been shown that the approximation ratios achievable by DSIC mechanisms and their non-incentive compatible counterparts exhibit a large asymptotic gap for some problems [Papadimitriou et al. 2008; Dobzinski 2011; Dughmi and Vondrák 2011; Dobzinski and Vondrak 2012].

Author's addresses: B. Lucier, Microsoft Research, brlucier@microsoft.com; A. Borodin, Department of Computer Science, University of Toronto, bor@cs.toronto.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0004-5411/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

Alternatively, one might study classes of “natural” allocation algorithms, that appear intuitive as auction allocation rules, with the hope that they have desirable incentive properties when implemented as mechanisms. As it turns out, for many combinatorial allocation problems, conceptually simple deterministic algorithms (e.g. greedy algorithms) meet or approach the best-known approximation factors subject to computational constraints [Lehmann et al. 1999; Mu’alem and Nisan 2008; Briest et al. 2005; Babaioff and Blumrosen 2004]. These natural methods tend to be computationally efficient and easy for bidders to understand, which are desirable properties in auctions. Unfortunately, such algorithms are not, in general, DSIC [Lehmann et al. 1999; Borodin and Lucier 2010]. Rather than abandoning these methods in favor of other, potentially more complex, mechanisms, we are pursuing an alternative approach. Namely, rather than striving for dominant strategy truthfulness, it may be acceptable for a system to admit strategic manipulation, so long as the designer’s objectives are met after such manipulation occurs. To this end, we explore the performance of mechanisms at equilibria of bidder behavior, given an appropriate model of beliefs. Broadly speaking, our motivating question is: When can an algorithm be implemented as a mechanism that achieves high social welfare at *every* equilibrium?¹ And how robust are the resulting mechanisms to variations of the equilibrium concept?

We demonstrate that for combinatorial allocation problems, any “greedy-like” approximation algorithm can be converted into a mechanism that achieves nearly the same approximation factor at every equilibrium of bidder behaviour. Our analysis is very general, and applies to a range of different equilibrium concepts, including pure and mixed Nash equilibria, Bayes-Nash (correlated) equilibria, no-regret equilibria, and iterated myopic best-response. We are thus able to decouple computational issues from incentives issues for this class of algorithms, as one can design a greedy algorithm without considering its economic implications, and then apply a straightforward pricing scheme in order to achieve good performance at equilibrium.

Performance of games at equilibrium has been studied extensively in the algorithmic game theory literature as the *price of anarchy (POA)* of a given game²: the ratio between the optimal outcome and the worst-case outcome at any equilibrium [Papadimitriou 2001]. Put into these terms, our goal is to convert an algorithm with approximation factor $c \geq 1$ into a mechanism whose price of anarchy is not much larger than c .

This paper is a synthesis and revision of results in [Lucier and Borodin 2010], [Lucier 2010], and results in the first author’s thesis [Lucier 2011]. The paper is organized as follows. The remainder of this section outlines our results and relates our work to recent papers in this area. Section 2 defines the necessary concepts and applications for our results. Section 3 introduces the concept of strongly loser-independence (generalizing the loser-independence concept from [Chekuri and Gamzu 2009]) which becomes the key property of greedy algorithms that we will exploit. Sections 4 and 5 analyze (respectively) price of anarchy results for first-price and critical-price mechanisms. In Sections 6 and 7 we consider solution concepts for repeated games: under regret minimization and best-response dynamics, respectively. Section 8 concludes with some open problems.

¹Dominant 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 [Christodoulou et al. 2008], Bayesian Nash equilibrium is not, strictly speaking, a relaxation of dominant strategy truthfulness. There exist truthful mechanisms whose approximation ratios are not preserved at all Nash equilibria, such as the famous Vickrey auction.

²For the purpose of this paper, we shall not consider cost minimization problems. We note that the price of anarchy concept was introduced in terms of cost minimization games but to the best of our knowledge the only price of anarchy results for mechanism induced games apply to maximization problems.

1.1. Our Results

The basic question of algorithmic mechanism design is this: when can computationally efficient algorithms be converted into mechanisms that preserve approximation bounds with respect to truthfulness or POA results? We address the price of anarchy considerations with respect to social welfare maximization for a broad class of allocation problems. In the full information and Bayesian setting, we study the price of anarchy for first price and critical price mechanisms derived from greedy algorithms. Roughly speaking (and in contrast to results regarding approximation and truthful mechanisms) we are able to show that there is often little or no loss from the approximation ratio of a greedy algorithm to the corresponding mechanism price of anarchy. We also study the long term behavior of the use of these mechanisms when used in repeated games.

One-Shot Auctions. We first consider one-shot auctions, in which the allocation problem is resolved only once. Following Christodoulou et al. [Christodoulou et al. 2008], we focus our attention on the standard (in economics) *incomplete information setting*, where the appropriate equilibrium concept is Bayes-Nash equilibrium. That is, we assume that agents' preferences are private, but drawn independently from commonly-known prior distributions, and that players apply strategies at equilibrium given this partial knowledge. We pose the question: can a given black-box approximation algorithm be converted into a mechanism that approximately preserves its approximation ratio at every Bayes-Nash Equilibrium? We show that for a broad class of greedy algorithms, the answer is *yes*.

Theorem (informal): Suppose \mathcal{A} is a greedy c -approximate allocation rule for a combinatorial allocation problem. Then the auction that uses \mathcal{A} to choose allocations, and uses a pay-your-bid payment scheme, has a Bayes-Nash Price of Anarchy of at most $c + O(c^2/e^c)$.

We also show that the small (and exponentially decreasing) loss in our price of anarchy bound is necessary, by giving an example (for every $c \geq 2$) where the resulting price of anarchy is at least $c + \Omega(\frac{c}{e^{4c}})$.

We note that the mechanisms we consider are all *prior-free*. Thus, as in the full-information case, while we assume the existence of type distributions in order to model rational agent behaviour, our mechanism need not be aware of these distributions. In the special case that each player's type distribution is a point mass, Bayes-Nash equilibrium reduces to standard Nash equilibrium. Our mechanisms therefore also preserve approximation ratios at every (mixed or pure) Nash equilibrium of the full information game. Our analysis also extends to the more general class of coarse correlated equilibria. For the case of pure Nash equilibrium, our price of anarchy bound improves to c .

As is standard, our bounds on the Bayesian price of anarchy will assume that agent types are distributed independently. However, we show that a weaker bound of $O(c)$ holds when agent types are drawn from an arbitrary distribution over the space of all type profiles. This result applies to greedy algorithms that are non-adaptive, as described in Section 2.4. Thus, even if agent types are arbitrarily correlated, our mechanisms yield performance at equilibrium asymptotically matching that of the underlying allocation algorithm.

A similar bound also applies to mechanisms that use the critical-price payment scheme, which is a natural extension of second-price payments in single-item auctions. Such a payment scheme charges each bidder the minimum bid at which he would have maintained his allocation. These bounds require a standard no-overbidding assumption.

tion, which is that agents are avoid bidding more than their value for any given subset of items.

Theorem (informal): Suppose \mathcal{A} is a greedy c -approximate allocation rule for a combinatorial allocation problem. Then, under the assumption that agents do not over-bid, the auction that uses \mathcal{A} to choose allocations, and uses a critical-price payment scheme, has a Bayes-Nash Price of Anarchy of at most $c + 1$.

We also show that the extra $+1$ term is necessary, by giving an example for every $c \geq 2$ in which the resulting price of anarchy is exactly $c + 1$. As with the first-price results, our bounds extend to coarse correlated equilibria, and a bound of $O(c)$ holds if agent valuations can be correlated. Furthermore, we show that a slight modification to the mechanism allows us to replace the no-overbidding assumption with the (conceptually weaker) assumption that bidders avoid weakly dominated strategies.

Repeated Auctions. Our bounds on efficiency at equilibrium do not explicitly model the manner by which agents reach equilibrium, or impose upon the agents any computational constraints whatsoever. We simply posit that equilibria (or approximate equilibria³) are predictive of the behavior of rational agents in high-stakes auctions. However, in settings where auctions are explicitly repeated, one might naturally model the dynamics under which bidder behavior evolves.

We will therefore also consider a *repeated-game* variant of combinatorial allocation problems, in which an auction problem is resolved multiple times with the same objects and bidders. Perhaps the most well-studied modern examples of repeated auctions are auctions for advertising spaces or slots [Edelman et al. 2005], but this model applies also to bandwidth auctions (such as the FCC spectrum auction), airline landing rights auctions [Cramton et al. 2005], etc. In these settings a mechanism for the (one-shot) auction problem corresponds to a repeated game to be played by the agents. Rather than view a repeated auction as an extensive-form game, we consider models of limited rationality that attempt to capture natural bidding behaviour. We study two such models: external regret minimization and asynchronous best-response dynamics.

In the first model, agents can play arbitrary sequences of strategies for the repeated auction, under the assumption that they obtain low regret relative to the best fixed strategy in hindsight. More precisely, for each bidder, the difference between the average utility obtained by the bidder and the average utility that would have been obtained by the best single declaration in hindsight must tend to 0 as the number of auction rounds increases. Under the assumption that bidders are able to minimize external regret, our goal is to design an auction mechanism that achieves an approximation to the optimal social welfare *on average* over sufficiently many rounds of the repeated auction. This is precisely the problem of designing a mechanism with bounded *price of total anarchy*, as introduced by Blum et al [Blum et al. 2008]. As observed by Blum and Mansour [Blum and Mansour 2007] and Roughgarden [Roughgarden 2009; 2012], in the full information setting, price of anarchy with respect to coarse correlated equilibrium is equal to the total price of anarchy. Hence, the bounds stated above for coarse correlated equilibrium apply also to the total price of anarchy. We further show that for the greedy mechanisms we consider, regret-minimizing strategies can be computed efficiently, assuming a natural representation of the bidders' valuation functions.

³All of our bounds on social efficiency degrade gracefully when agents apply strategies in approximate equilibrium. Namely, whenever we convert a c approximate allocation algorithm into a mechanism achieving say $f(c)$ price of anarchy, the same proof shows that the mechanism achieves at least (and often better) an $f(c + \gamma)$ approximation at every $(1 + \gamma)$ approximate equilibrium. Notably, if $c \geq 1 + \gamma$, the c -approximation for pure equilibria of the first price mechanism remains a c approximation at every approximate equilibrium.

Theorem (informal): Suppose \mathcal{A} is a greedy c -approximate allocation rule for a combinatorial allocation problem. Then the auction that uses \mathcal{A} to choose allocations, and charges critical payments, achieves a $(c + 1)$ approximation to the optimal welfare when agents apply regret-minimizing strategies. Moreover, a regret-minimizing strategy can be implemented in time polynomial in the number of XOR bids⁴ used to represent an agent's valuation.

In the second model, we assume that agents choose strategies that are myopic best-responses to the current strategies of the other agents. We model this behaviour as follows: on each auction round, an agent is chosen uniformly at random, and that agent is given the opportunity to change his strategy to the current myopic best-response. As in the regret-minimization model, our goal is to design auction mechanisms that achieve approximations to the best possible social welfare *on average* over sufficiently many auction rounds, with high probability over the random choices of bidders to update. This is the concept of the *price of (myopic) sinking*, as introduced by Goemans et al [Goemans et al. 2005].

We conjecture that any greedy c -approximate allocation rule can be implemented as a mechanism with price of sinking $O(c)$. As partial progress toward this conjecture, we design mechanisms tailored to two particular combinatorial allocation problems: the unrestricted combinatorial auction problem and the cardinality-restricted combinatorial auction. Each mechanism has price of sinking $O(c)$, where c is the approximation factor of the best-known algorithm.

We recall that one method for bounding the price of sinking is to prove that there exists an equilibrium state that is reachable from any declaration profile by some polynomial-length sequence of best-response steps. This would imply that an equilibrium state would be reached with high probability after exponentially many steps. We do not take this approach, but rather prove that the average social welfare obtained after a *polynomial* number of steps will approximate the optimal welfare with high probability.

1.2. Related Work

The seminal paper in algorithmic game theory and more specifically algorithmic mechanism design is that of Nisan and Ronen [Nisan and Ronen 1999]. The basic issue introduced in [Nisan and Ronen 1999] is to reconcile the competing demands for revenue and social welfare optimization with the need for computational efficiency in the context of self interested (i.e. selfish) agents. The two most studied solution concepts in algorithmic game theory are truthfulness (i.e. incentive compatibility) and behavior at (all) equilibria (i.e. the price of anarchy (POA) concept). Initial POA results for games were first introduced to algorithmic game theory in the seminal papers by Papadimitriou [Papadimitriou 2001] and Roughgarden and Tardos [Roughgarden and Tardos 2000]. Christodoulou et al. [Christodoulou et al. 2008] initiated the study of the price of the price of anarchy in the Bayesian setting. Whereas the emphasis of algorithmic mechanism design has been to consider the approximations achievable by truthful mechanisms, to the best of our knowledge, our conference paper [Lucier and Borodin 2010] was the first to consider this constructive aspect of mechanism design and price of anarchy.

Since the initial conference version of this work there has been significant progress on the understanding of the price of anarchy of mechanisms in various auction settings. Some examples include the Generalized Second Price auction for sponsored

⁴Equivalently, the minimum number of (subset, value) pairs (S_i, w_i) needed so that valuation v satisfies $v(T) = \max_{S_i} v(S_i \cup T) - w_i$.

search ads [Caragiannis et al. 2012], simultaneous single-item auctions [Christodoulou et al. 2008; Hassidim et al. 2011; Bhawalkar and Roughgarden 2011; Feldman et al. 2013], and multi-unit auctions [Markakis and Telelis 2012; de Keijzer et al. 2013]. A framework unifying much of this work was proposed by Syrgkanis and Tardos [Syrgkanis and Tardos 2013].

Chekuri and Gamzu [Chekuri and Gamzu 2009] defined “loser-independent algorithms”, and in the conference version of our paper [Lucier and Borodin 2010] we argued that the basic property of greedy algorithms that we were exploiting was a multi-parameter version of loser-independence. In the first author’s thesis [Lucier 2011], a strengthening of loser-independence, called strong loser-independence, was introduced to simplify the proofs and that will be the basic property of greedy algorithms we will use in this paper. Loser-independence is conceptually related to the concept of *smoothness*, which was introduced by Roughgarden [Roughgarden 2009] as a general way to derive price of anarchy results for one shot and repeated games (without reference to mechanisms that derive games). Loser-independence has been shown to be different from this original notion of smoothness [Lucier 2011]. However, alternative notions of smoothness defined by Lucier and Paes Leme [Lucier and Leme 2011] and Syrgkanis and Tardos [Syrgkanis and Tardos 2013] can also be used to derive results similar to our results. In particular, Syrgkanis and Tardos use their smoothness condition to derive many price of anarchy results for allocation mechanisms, including those derived from greedy c -approximation algorithms. Their result for the (non correlated) mixed Bayesian and coarse correlated equilibrium improved upon our conference results: as in our current paper, they show that the resulting price of anarchy approaches c with a term exponentially decreasing in c . In particular, they show that the price of anarchy is never worse than $c + .58$. As we will show in Section 3.1, our application of strong loser-independence can be interpreted as a proof of smoothness.

2. PRELIMINARIES

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 $y \geq 0$ such that for all $T \subseteq M$, $v(T) = y$ if $S \subseteq T$ and 0 otherwise. 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. The goal in our social welfare maximization problems is to choose an allocation for each agent in order to maximize the sum of agent values.

A *combinatorial allocation problem* is defined by a set of *feasible allocations*, which is the set of permitted allocation profiles. We further assume in combinatorial allocation problems that this feasibility constraint is *separable*, meaning that if \mathbf{x} is feasible then $(\emptyset, \mathbf{x}_{-i})$ is also feasible⁵ for all i . Note that separability is a weaker assumption than the standard downward-closure property of packing problems, which would stipulate that if \mathbf{x} is feasible then (y_i, \mathbf{x}_{-i}) is also feasible for all $y_i \subseteq x_i$. 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 . An allocation rule is *component-wise monotone* if it satisfies the

⁵We note that the combinatorial public projects problem (CPPP) [Papadimitriou et al. 2008] is not separable.

following property for every agent i :

If $v_i(S) < \tilde{v}_i(S)$, $v_i(T) = \tilde{v}_i(T) \forall T \neq S$, and $\mathcal{A}_i(v_i, \mathbf{v}_{-i}) = S$, then $\mathcal{A}_i(\tilde{v}_i, \mathbf{v}_{-i}) = S$

We will tend to write \mathcal{A} for both an allocation rule and an algorithm that implements it. We will sometimes abuse notation and use \mathbf{x} for an allocation rule, rather than a specific allocation.

Each agent $i \in [n]$ has a private valuation function v_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{v} , is $SW(\mathbf{x}, \mathbf{v}) = \sum_i v_i(x_i)$. We write $SW_{opt}(\mathbf{v})$ for $\max_{\mathbf{x}} \{SW(\mathbf{x}, \mathbf{v})\}$ and say that algorithm \mathcal{A} is a c approximation algorithm⁶ if $SW(\mathcal{A}(\mathbf{v}), \mathbf{v}) \geq \frac{1}{c} SW_{opt}(\mathbf{v})$ for all \mathbf{v} .

A type profile \mathbf{v} and an allocation rule \mathcal{A} for a combinatorial allocation problem define *critical values*, $\theta_i(S, \mathbf{v}_{-i})$, for any agent i and set $S \subseteq M$. The value $\theta_i(S, \mathbf{v}_{-i})$ is the minimum amount that agent i could bid on set S and still win S , assuming the other agents have profile \mathbf{v}_{-i} . That is, $\theta_i(S, \mathbf{v}_{-i}) = \inf\{z : x_i(z, \mathbf{v}_{-i}) = S\}$. We note that this definition of critical values holds even if it is not the case that increasing one's value for a set necessarily increases the probability of obtaining that set. However, most of the mechanisms we consider in this work do satisfy this monotonicity property, which motivates the terminology of a critical price.

2.2. Mechanisms

A *direct revelation mechanism* $\mathcal{M}(\mathcal{A}, P)$ is composed of an allocation rule \mathcal{A} and a payment rule P that assigns a vector of n payments to each declared valuation profile. 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, we do not assume that \mathbf{d} is equal to \mathbf{v} . We will write $SW(\mathbf{d})$ for $SW(\mathcal{A}(\mathbf{d}), \mathbf{v})$ when the allocation rule and type profile are clear from context.

We will be concerned with two different payment rules, *first price* and *critical price*. In a first price mechanism, an agent is charged their declared bid $d_i(S)$ for any allocated set S . For notational convenience, we let $\mathcal{M}_1(\mathcal{A})$ denote the mechanism using allocation rule \mathcal{A} and the first price payment rule. In the critical price payment rule, an agent is charged his critical value $\theta_i(S, \mathbf{d}_{-i})$ for any allocated set S . We will let $\mathcal{M}_2(\mathcal{A})$ denote the mechanism using allocation rule \mathcal{A} and the critical price payment rule.

2.3. Equilibria of One-shot Auctions

The utility of agent i in mechanism $\mathcal{M} = (\mathcal{A}, P)$, given declaration profile \mathbf{d} and type profile \mathbf{v} , is $u_i^{v_i}(\mathbf{d}) = v_i(\mathcal{A}_i(\mathbf{d})) - P_i(\mathbf{d})$. We will often omit the dependence on v_i when it is clear from context, and write simply $u_i(\mathbf{d})$. We say that declaration d_i weakly dominates d'_i if for all \mathbf{d}_{-i} , $u_i(d_i, \mathbf{d}_{-i}) \geq u_i(d'_i, \mathbf{d}_{-i})$, and that there exist at least one \mathbf{d}_{-i} for which the inequality is strict.

We consider a Bayesian setting in which the true types of the agents are not fixed, but are rather drawn from a known probability distribution \mathbf{F} over the set of valuation profiles. We first assume that $\mathbf{F} = F_1 \times \dots \times F_n$ is the product of independent distributions, where $F_i(v_i)$ is the probability that agent i has type v_i . (Later we will also consider correlated distributions over type profiles.) We write $SW_{opt}(\mathbf{F})$ for $\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[SW_{opt}(\mathbf{v})]$.

A *bidding strategy* for agent i is a function b_i that maps a type v_i to a distribution over declarations for agent i . We think of $b_i(v_i)$ as the (randomized) bidding strategy

⁶Our convention will be to have approximation ratios $c \geq 1$.

employed by agent i given that his true type is v_i . We will abuse notation slightly and also write $b_i(v_i)$ for the random variable representing a declaration chosen from the corresponding distribution. We write $\mathbf{b}(\mathbf{v}) = b_1(v_1) \times \dots \times b_n(v_n)$ for the (distribution over) declaration profiles resulting from applying the bid functions in \mathbf{b} to type profile \mathbf{v} . The strategy profile \mathbf{b} forms a *(mixed) Bayesian Nash Equilibrium* (BNE) if, for every $i \in [n]$ and every v_i in the support of F_i , agent i maximizes his expected utility by making a declaration drawn from distribution $b_i(v_i)$. That is, for each agent i , each possible type v_i , and every distribution ω_i' over declarations,

$$\mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} [u_i(\mathbf{b}(\mathbf{v}))] \geq \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}, d_i \sim \omega_i'} [u_i(d_i, \mathbf{b}_{-i}(\mathbf{v}_{-i}))].$$

For a mechanism $\mathcal{M} = (\mathcal{A}, P)$, we will write $SW_{\mathcal{M}}(\mathbf{F}, \mathbf{b})$ to mean $\mathbf{E}_{\mathbf{v} \sim \mathbf{F}} [\sum_i v_i(\mathcal{A}_i(\mathbf{b}(\mathbf{v})))]$, the expected social welfare given type distribution \mathbf{F} and strategy profile \mathbf{b} .

The *(mixed) Bayesian price of anarchy* (BPoA) of mechanism \mathcal{M} is defined as

$$\sup_{\mathbf{F}, \mathbf{b}} \frac{SW_{opt}(\mathbf{F})}{SW_{\mathcal{M}}(\mathbf{F}, \mathbf{b})}$$

where the supremum is over all type distributions \mathbf{F} and mixed BNE \mathbf{b} for \mathbf{F} . In other words, the BPoA of \mathcal{M} is the worst-case ratio between the expected welfare at BNE and the expected optimal welfare.

We can further extend the definition of BNE to allow a correlated distribution over type profiles. The definition for correlated BNE and correlated Bayesian price of anarchy is then the same as the above definitions, where we would no longer assume that F is a product of independent distributions.

Returning to the case in which F is a product distribution, a number of special cases deserve mention. When all type distributions are point masses (i.e., each agent's type is determined), a BNE is referred to as a *(mixed) Nash Equilibrium* (NE). The Price of Anarchy (PoA) of a mechanism \mathcal{M} is defined analogously to the BPoA, but with respect to fixed type profiles and mixed Nash equilibria. It follows that the BPoA is always at least the PoA for a given mechanism. A BNE (or NE) is called *pure* if its constituent bidding strategies are deterministic. In general a pure Nash equilibrium may not exist for a given mechanism and type profile; see Appendix A.

One can generalize mixed NE by relaxing the assumption that the declaration distributions are independent. That is, one might allow $\mathbf{b}(\mathbf{v})$ to be an arbitrary distribution over declarations, rather than a product distribution. A distribution ω over declaration profiles is a *coarse correlated equilibrium* (CCE) for type profile \mathbf{v} if, for all i and all declaration distributions ω_i' ,

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

Note that when the agent declaration distributions are independent, CCE is equivalent to mixed NE. We define the analogous price of anarchy concepts; it follows that the pure price of anarchy is at most the mixed price of anarchy which in turn is at most the coarse correlated price of anarchy.

2.4. 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 priority framework in Borodin, Nielsen and Rackoff [Borodin et al. 2002] and the monotone greedy algorithms of Mu'alem and Nisan [Mu'alem and Nisan 2008], extended to be adaptive as in Borodin and Lucier [Borodin and Lucier 2010]. We begin with some definitions. A *partial allocation profile*

Priority Algorithm

Input: Declaration profile $\mathbf{d} = d_1, \dots, d_n$.

1. Fix a monotone priority function r . Let $N = \emptyset$.
2. Repeat until $N = [n]$:
3. Choose $i \notin N$ and feasible allocation $S \subseteq M$ for i given N that maximizes $r(i, S, d_i(S))$
4. Set $x_i = S$; add player i to N
5. return $\mathbf{x} = (x_1, \dots, x_n)$

Fig. 1. The framework for a non-adaptive priority algorithm.

for agents $N \subseteq [n]$ is an allocation profile \mathbf{x} with $x_i = \emptyset$ for all $i \notin N$. A partial allocation profile is *feasible* if there is some feasible allocation profile that extends it. Given a partial allocation profile \mathbf{x} for subset N , some $i \notin N$, and allocation $y \subseteq M$, we say y is a *feasible allocation for i given N* if the partial allocation remains feasible when we set $x_i = y$.

A *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 say that r is monotone if it is non-decreasing in v and monotone non-increasing in S with respect to set inclusion.

We consider two types of greedy allocation algorithms. A *non-adaptive greedy allocation algorithm* \mathcal{A} is an allocation algorithm as defined in Figure 1. We say that \mathcal{A} is monotone when the priority function r is monotone. We assume that ties in step 3 are broken in an arbitrary but fixed manner (i.e. we assume that the priority function is a 1-1 function inducing a total ordering).

A non-adaptive algorithm fixes a single priority function that is used throughout its execution. By contrast, an *adaptive greedy allocation algorithm* can change its priority function on each iteration, depending on the partial allocation formed on the previous iterations.

2.5. Applications

We now describe some applications of greedy algorithms for particular combinatorial allocation problems.

Combinatorial Auctions. The general combinatorial auction problem is defined by the feasibility constraint that no two allocations can intersect. Lehmann et al [Lehmann et al. 1999] show that the (non-adaptive) greedy allocation rule with $r(i, S, v) = \frac{v}{\sqrt{|S|}}$ achieves a $\sqrt{2m}$ approximation ratio for CAs.

Cardinality-restricted Combinatorial Auctions. In the special case that players' desires are restricted to sets of size at most k , the non-adaptive greedy algorithm with $r(i, S, v) = v$ is k -approximate assuming single-minded agents. This translates to a $(k + 1)$ approximate algorithm for general (i.e. multi-minded) agents.

Multiple-Demand Unsplittable Flow Problem. 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 $B > 1$ [Briest et al. 2005].

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 [Babaioff and Blumrosen 2004]. Alternative greedy algorithms yield better approximation ratios for special cases, such as rectangles.

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 greedy algorithm that orders bids by value obtains a 3-approximation, and is symmetric with respect to agents and objects.

Unlike the previous examples, for the case of single-minded bidders, there is an optimal dynamic programming algorithm that runs in time $O(n^7)$ [Baptiste 1999]. Since this algorithm solves the problem optimally, it is incentive compatible. In this case, the resulting price of anarchy for the greedy algorithm is appealing primarily due to its linear runtime and simple allocation rule.

3. STRONG LOSER-INDEPENDENCE

Chekuri and Gamzu [Chekuri and Gamzu 2009] introduced a property known as loser-independence for combinatorial allocation algorithms in single-parameter domains. They define an algorithm for a combinatorial allocation problem to be *loser-independent* if, whenever $\mathcal{A}_i(d_i, \mathbf{d}_{-i}) = \mathcal{A}_i(d_i', \mathbf{d}_{-i}) = \emptyset$ for some i , \mathbf{d}_{-i} , d_i , and d_i' , then it must be that $\mathcal{A}(d_i, \mathbf{d}_{-i}) = \mathcal{A}(d_i', \mathbf{d}_{-i})$. That is, if a “losing” agent (i.e. an agent who is allocated no items) modifies his declaration in such a way that he still receives no items, this cannot affect the outcome of algorithm \mathcal{A} . Note that loser-independence is a condition on declaration profiles, rather than on bidding functions, since the loser-independence notion is purely algorithmic and is not a condition on equilibria. In our results we will make use of a stronger property of greedy algorithms, which we call strong loser-independence.

DEFINITION 3.1. *An allocation rule \mathcal{A} is strongly loser-independent if, whenever \mathbf{d} and \mathbf{d}' satisfy $\mathcal{A}(\mathbf{d}) \neq \mathcal{A}(\mathbf{d}')$, there exists an agent i and set $S \neq \emptyset$ such that $d_i(S) \neq d_i'(S)$ and either $\mathcal{A}_i(\mathbf{d}) = S$ or $\mathcal{A}_i(d_i', \mathbf{d}_{-i}) = S$.*

Roughly speaking, if \mathcal{A} is a strongly loser-independent algorithm, then whenever a valuation profile changes from \mathbf{d} to \mathbf{d}' via modifications to “losing bids” (i.e. an agent i ’s declared value for sets that are not allocated to him, when others bid according to \mathbf{d}_{-i}), algorithm \mathcal{A} will return the same outcome on inputs \mathbf{d} and \mathbf{d}' . We note that our definition requires that either $\mathcal{A}_i(\mathbf{d}) = S$ or $\mathcal{A}_i(d_i', \mathbf{d}_{-i}) = S$, rather than $\mathcal{A}_i(\mathbf{d}') = S$. The intuition is that we think of “losing bids” as being losers with respect to the original declaration profile \mathbf{d} .

The property of strong loser-independence strengthens the definition of loser-independence due to Chekuri and Gamzu in two ways. First, we extend from single-parameter settings to multiple-parameter settings by considering losing bids rather

than losing agents. Second, we require that the algorithm outcome be unaffected if multiple agents simultaneously modify losing bids.

It is clear from the definitions that all strongly loser-independent algorithms are loser-independent (i.e. by considering the case when \mathbf{d} and \mathbf{d}' differ only on the declaration of a single agent). However, not all loser-independent algorithms are strongly loser-independent, even in single-minded domains. For example, consider the combinatorial auction problem and suppose that \mathcal{A} is an algorithm that optimizes social welfare exactly and breaks ties consistently. Then \mathcal{A} is loser-independent, since a losing agent's bid does not affect the optimal allocation. However, \mathcal{A} is not strongly loser-independent, as the following instance shows. Consider an auction of two items $\{a, b\}$ to three bidders. If the (single-minded) bidder declarations are $d_1(\{a, b\}) = 10$, $d_2(\{a\}) = 3$, and $d_3(\{b\}) = 3$, then the outcome is that agent 1 wins his desired set. On the other hand, if the bidder declarations are given by \mathbf{d}' where $d_1'(\{a, b\}) = 10$, $d_2'(\{a\}) = 6$, and $d_3'(\{b\}) = 6$, then the outcome changes: agents 2 and 3 win their desired sets. However, \mathbf{d} and \mathbf{d}' do not differ in their declarations for any sets allocated by $\mathcal{A}(\mathbf{d})$, $\mathcal{A}(d_1', \mathbf{d}_{-1})$, $\mathcal{A}(d_2', \mathbf{d}_{-2})$, or $\mathcal{A}(d_3', \mathbf{d}_{-3})$, as agent 1 wins his desired set in each of these four cases. This contradicts the definition of strongly loser-independence.

As we now show, all greedy algorithms satisfy the strong loser-independence property.

LEMMA 3.2. *Every (monotone) adaptive greedy algorithm is (component-wise monotone) and strongly loser-independent.*

PROOF. The monotonicity property follows immediately when the priority function in the greedy algorithm is a monotone function.

Let \mathcal{A} be an adaptive greedy allocation rule, and choose any \mathbf{d} and \mathbf{d}' such that $\mathcal{A}(\mathbf{d}) \neq \mathcal{A}(\mathbf{d}')$. We will show that there exists some i and S such that $d_i(S) \neq d_i'(S)$ and either $\mathcal{A}_i(\mathbf{d}) = S$ or $\mathcal{A}_i(d_i', \mathbf{d}_{-i}) = S$.

Recall the definition of an adaptive greedy algorithm, and consider the iterations of \mathcal{A} on inputs \mathbf{d} and \mathbf{d}' . Let k be the first iteration in which the allocation of \mathcal{A} differs on these two inputs. Suppose that \mathcal{A} allocates set U to agent ℓ on iteration k when the input is \mathbf{d} , and allocates T to agent j on iteration k when the input is \mathbf{d}' .

For each iteration $q < k$, write i_q for the agent allocated to by \mathcal{A} (on either input profile) and S_q for the set allocated to i_q . Note that if $d_{i_q}(S_q) \neq d_{i_q}'(S_q)$ for any $q < k$ then we have the desired result with $i = i_q$ and $S = S_q$. We can therefore assume that $d_{i_q}(S_q) = d_{i_q}'(S_q)$ for all $q < k$. This implies that the bids resolved by \mathcal{A} are identical on all iterations preceding k on inputs \mathbf{d} and \mathbf{d}' , and therefore the values of ranking functions used in each iteration up to k must be identical for inputs \mathbf{d} and \mathbf{d}' . Write r_q for the ranking function used in iteration q for each $q \leq k$. Thus, since the allocation on iteration k changed from choosing set U for agent ℓ to choosing set T for agent j , it must be that either $r_k(\ell, U, d_\ell(U)) \neq r_k(\ell, U, d_\ell'(U))$ or $r_k(j, T, d_j(T)) \neq r_k(j, T, d_j'(T))$. This implies that either $d_\ell(U) \neq d_\ell'(U)$ or $d_j(T) \neq d_j'(T)$.

If $d_\ell(U) \neq d_\ell'(U)$ then we have the desired result with $i = \ell$ and $S = U$, since $\mathcal{A}_\ell(\mathbf{d}) = U$. We can therefore assume that $d_\ell(U) = d_\ell'(U)$ and $d_j(T) \neq d_j'(T)$. Consider now the behaviour of algorithm \mathcal{A} on input (d_j', \mathbf{d}_{-j}) . We claim that $\mathcal{A}_j(d_j', \mathbf{d}_{-j}) = T$. Note that this implies the desired result with $i = j$ and $S = T$. To prove the claim, recall that $d_{i_q}(S_q) = d_{i_q}'(S_q)$ for all $q < k$. Thus, for each $q < k$ and each feasible set S that could be allocated to agent j on iteration q ,

$$r_q(i_q, S_q, d_{i_q}(S_q)) = r_q(i_q, S_q, d_{i_q}'(S_q)) > r_q(j, S, d_j'(S))$$

since \mathcal{A} allocates S_q to i_q on input \mathbf{d}' . We conclude that, on input (d_j', \mathbf{d}_{-j}) , \mathcal{A} allocates S_q to agent i_q on each iteration $q < k$. On iteration k , we have $r_k(j, T, d_j'(T)) >$

$r_k(j, T', d_j'(T'))$ for any feasible $T' \neq T$ (since \mathcal{A} allocates T to j on input d') and

$$r_k(j, T, d_j'(T)) > r_k(\ell, U, d_\ell'(U)) = r_k(\ell, U, d_\ell(U)) \geq r_k(i, S, d_i(S))$$

for any feasible $i \neq j$, due to our assumption that $d_\ell'(U) = d_\ell(U)$ and the fact that \mathcal{A} allocates U to ℓ on iteration k for input d . We therefore conclude that $\mathcal{A}_j(d_j', d_{-i}) = T$ as required. \square

We next explore an implication of a strongly loser-independent algorithm \mathcal{A} being a worst-case c -approximation. 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 values of the optimal allocation profile.

LEMMA 3.3. *If \mathcal{A} is a c -approximate strongly loser-independent algorithm, then for any type profile \mathbf{v} and allocation profile \mathbf{y} , $\sum_{i \in [n]} v_i(\mathcal{A}_i(\mathbf{v})) \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(y_i, \mathbf{v}_{-i})$.*

PROOF. Choose any $\epsilon > 0$. For all i , let v_i' be the single-minded declaration for set y_i at value $\theta_i(y_i, \mathbf{v}_{-i}) - \epsilon$. Let v_i^* be the pointwise maximum of v_i' and v_i . That is, for all $S \subseteq M$, $v_i^*(S) = \max\{v_i(S), v_i'(S)\}$. By definition of critical prices, we have that $\mathcal{A}_i(v_i^*, \mathbf{v}_{-i}) = \mathcal{A}_i(\mathbf{v})$ for all i , and furthermore $v_i^*(\mathcal{A}_i(\mathbf{v})) = v_i(\mathcal{A}_i(\mathbf{v}))$. Since \mathcal{A} is strongly loser-independent, we must therefore have $\mathcal{A}(\mathbf{v}) = \mathcal{A}(\mathbf{v}^*)$. Since \mathcal{A} is a c -approximation, we conclude that $SW(\mathbf{x}(\mathbf{v}), \mathbf{v}) = SW(\mathbf{x}(\mathbf{v}^*), \mathbf{v}^*) \geq \frac{1}{c} SW(\mathbf{y}, \mathbf{v}^*) \geq \frac{1}{c} \sum_{i \in [n]} \theta_i(y_i, \mathbf{v}_{-i}) - n\epsilon$. The result follows by taking the limit as $\epsilon \rightarrow 0$. \square

For brevity, for the remainder of this paper we will say “monotone strongly loser-independent” to mean both strongly loser-independent and component-wise monotone⁷.

3.1. Applying Strong Loser Independence

Strong loser-independence is a strictly algorithmic concept, devoid of game theoretic considerations. Our general approach will be to derive price of anarchy results for any mechanism that uses a strongly loser-independent c -approximation \mathcal{A} as its allocation algorithm. To do so, we will be using Lemma 3.3 in conjunction with the assumption that a given bid profile is an equilibrium.

At a high level, our argument will be as follows. For each pricing rule and equilibrium concept, equilibrium will imply an inequality of the form

$$v_i(y_i) \leq \lambda \cdot \theta_i(y_i, \mathbf{d}_{-i}) + \mu \cdot v_i(x_i(\mathbf{d}))$$

where \mathbf{y} is an optimal allocation. (For Bayesian equilibrium, these terms will be expectations.) This allows us to charge the optimal gain for each agent to its critical value and its welfare from the algorithm. We then exploit Lemma 3.3 to convert this bound into a relationship between the optimal welfare and the welfare at equilibrium. To make this more specific, in our pure Nash equilibrium result for a first price mechanism (Theorem 4.3), we show the following (somewhat stronger) inequality:

$$v_i(y_i) \leq \theta_i(y_i, \mathbf{d}_{-i}) + v_i(x_i(\mathbf{d})) - d_i(x_i(\mathbf{d}))$$

⁷For pure Nash equilibrium price of anarchy results, if we assume no over-bidding, we do not need monotonicity but it is necessary for all of our other results.

It will then follow (ignoring an $n\epsilon$ term which disappears as $\epsilon \rightarrow 0$) that:

$$\begin{aligned} \sum_i v_i(y_i) &\leq \sum_i \theta_i(y_i, \mathbf{d}_{-i}) + \sum_i v_i(x_i(\mathbf{d})) - \sum_i d_i(x_i(\mathbf{d})) \\ &\leq (c-1) \sum_i d_i(x_i(\mathbf{d})) + \sum_i v_i(x_i(\mathbf{d})) \end{aligned}$$

In other words, the high-level approach is to charge an agent's welfare in the optimal outcome against his welfare at equilibrium plus the welfare of other "price-setting" agents. This approach is similar to the smoothness argument as formulated by Syrgkanis and Tardos [Syrgkanis and Tardos 2013]. However, there is a difference in our approaches. The smoothness condition in [Syrgkanis and Tardos 2013] is tailored to allocation mechanisms and asserts the existence of some d_i (for each player i) satisfying such an inequality whereas we are assuming that \mathbf{d} is an equilibrium. The benefit of their immediate reduction to smoothness is that their price of anarchy results for pure equilibria carry over immediately to Bayesian price of anarchy. However, this prohibits establishing certain tight bounds; for example, in the first price mechanism we show that the pure price of anarchy is c , which cannot be achieved via smoothness since this bound does not hold for the Bayesian price of anarchy.

4. FIRST-PRICE MECHANISMS

In this section we analyze greedy algorithms paired with a first price payment scheme. More precisely (with the exception of results relating to correlated Bayesian equilibrium where we will consider more specific greedy allocations), given a strongly loser-independent algorithm \mathcal{A} , we will be studying the performance of the first-price mechanism $\mathcal{M}_1(\mathcal{A})$ at equilibrium.

Our first step will be to show that a utility-maximizing declaration of an agent never involves overbidding on a set that he may possibly be allocated. This will imply that agents do not employ overbidding strategies at equilibrium. It may appear at first glance that any strategy that recommends overbidding on sets is obviously dominated for any allocation algorithm, since winning any bid larger than one's true value leads to negative utility. However, we must also show that an agent cannot find it advantageous to overbid on some set S in order to affect his probability of winning some other set T . We will demonstrate that such situations cannot occur when allocations are chosen by a strongly loser independent algorithm.

For a type v_i and a declaration d_i , we will write \bar{d}_i for the declaration defined as $\bar{d}_i(S) = \min\{v_i(S), d_i(S)\}$. That is, \bar{d}_i agrees with d_i , except that the declared value of each set can be at most the true value for that set. Note that $\bar{d}_i = d_i$ precisely if d_i does not overbid on any set.

We now show that any declaration d_i that overbids on a set that could potentially be won is weakly dominated by strategy \bar{d}_i .

LEMMA 4.1. *For any monotone strongly loser independent allocation rule \mathcal{A} , valuation v_i , and declaration profile \mathbf{d} , we have $u_i(\mathbf{d}) \leq u_i(\bar{d}_i, \mathbf{d}_{-i})$. Moreover, the inequality is strict when $d_i(\mathcal{A}(\mathbf{d})) > v_i(\mathcal{A}(\mathbf{d}))$.*

PROOF. Let $S = \mathcal{A}_i(\mathbf{d})$. Suppose first that $d_i(S) > v_i(S)$. Then $u_i(\mathbf{d}) = v_i(S) - d_i(S) < 0$. Since $v_i(T) - \bar{d}_i(T) \geq 0$ for every set T , this implies that $u_i(\bar{d}_i, \mathbf{d}_{-i}) > u_i(\mathbf{d})$, as required.

Next suppose that $d_i(S) \leq v_i(S)$, so that $\bar{d}_i(S) = d_i(S)$. We claim that $\mathcal{A}_i(\bar{d}_i, \mathbf{d}_{-i}) = S$. Suppose not, for contradiction. Then we can construct a sequence of declarations (d^1, d^2, \dots, d^k) , with $d^1 = d_i$ and $d^k = \bar{d}_i$, such that adjacent declarations differ only on a single set and declared values only decrease. Suppose j is minimal such that

$\mathcal{A}_i(d^j, \mathbf{d}_{-i}) \neq S$; such a $j > 1$ must exist since, by assumption, $\mathcal{A}_i(\bar{d}_i, \mathbf{d}_{-i}) \neq S$. Then (a) d^{j-1} and d^j differ only on the value assigned to some set T , (b) $d^{j-1}(T) > d_i^j(T)$, (c) $\mathcal{A}_i(d^{j-1}, \mathbf{d}_{-i}) = S$, and (d) $\mathcal{A}_i(d^j, \mathbf{d}_{-i}) \neq S$. Strong loser-independence then implies that $\mathcal{A}_i(d^j, \mathbf{d}_{-i}) = T$. However, the fact that $d^{j-1}(T) > d^j(T)$ then contradicts the component-wise monotonicity of \mathcal{A} .

We conclude by contradiction that $\mathcal{A}_i(\bar{d}_i, \mathbf{d}_{-i}) = S$. Since S is also $\mathcal{A}(\mathbf{d})$, we have $u_i(\mathbf{d}) = v_i(S) - d_i(S) = u_i(\bar{d}_i, \mathbf{d}_{-i})$ as required.

□

An immediate corollary is that if b_i is a bidding strategy, and there exists a type v_i and set S such that $(b_i(v_i))(S) > v_i(S)$, then b_i is weakly dominated by the strategy \bar{b}_i . Moreover, \bar{b}_i is strictly better, in terms of utility, under any distribution of declarations in which agent i wins set S with positive probability. We conclude that at any BNE of mechanism $\mathcal{M}_1(\mathcal{A})$, no player will overbid on a set that he wins with positive probability.

COROLLARY 4.2. *For any monotone strongly loser independent allocation rule \mathcal{A} , BNE \mathbf{b} , type v_i , and set S , if $\Pr_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[\mathcal{A}_i(\mathbf{b}(\mathbf{v})) = S] > 0$ then $(b_i(v_i))(S) \leq v_i(S)$.*

4.1. Pure Nash Equilibria

We are now ready to bound the price of anarchy of $\mathcal{M}_1(\mathcal{A})$. We begin with a result for pure Nash equilibria, rather than the fully general BNE case.

THEOREM 4.3. *Suppose \mathcal{A} is a c -approximate monotone strongly loser independent allocation rule for a combinatorial allocation problem. Then the price of anarchy of $\mathcal{M}_1(\mathcal{A})$ is at most c .*

PROOF. Fix type profile \mathbf{v} and suppose that \mathbf{b} forms a pure Nash equilibrium. Since the Nash equilibrium is pure, we will write $\mathbf{d} = \mathbf{b}(\mathbf{v})$ for notational convenience. Let \mathbf{y} be an optimal allocation for \mathbf{v} , and let $\mathbf{x}(\cdot)$ denote the allocation rule for \mathcal{A} . Lemma 3.3 implies

$$\sum_i d_i(x_i(\mathbf{d})) \geq \frac{1}{c} \sum_i \theta_i(y_i, \mathbf{d}_{-i}). \quad (2)$$

Choose arbitrarily small $\epsilon > 0$ and let d_i' be the single-minded declaration for set y_i at value $\theta_i(y_i, \mathbf{d}_{-i}) + \epsilon$. Then $x_i(d_i', \mathbf{d}_{-i}) = y_i$ (from the definition of critical values) and hence $u_i(d_i', \mathbf{d}_{-i}) = v_i(y_i) - \theta_i(y_i, \mathbf{d}_{-i}) - \epsilon$. Since \mathbf{d} is a Nash equilibrium, it must be that

$$\begin{aligned} v_i(y_i) - \theta_i(y_i, \mathbf{d}_{-i}) - \epsilon &= u_i(d_i', \mathbf{d}_{-i}) \\ &\leq u_i(d_i, \mathbf{d}_{-i}) \\ &= v_i(x_i(\mathbf{d})) - d_i(x_i(\mathbf{d})). \end{aligned}$$

Summing over all i and applying (2) and Corollary 4.2 we have

$$\begin{aligned} \sum_i v_i(y_i) &\leq \sum_i \theta_i(y_i, \mathbf{d}_{-i}) - \sum_i d_i(x_i(\mathbf{d})) + \sum_i v_i(x_i(\mathbf{d})) + n\epsilon \\ &\leq (c-1) \sum_i d_i(x_i(\mathbf{d})) + \sum_i v_i(x_i(\mathbf{d})) + n\epsilon \\ &\leq c \sum_i v_i(x_i(\mathbf{d})) + n\epsilon \end{aligned}$$

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

$$\begin{aligned} SW(\mathbf{x}(\mathbf{d}), \mathbf{v}) &= \sum_i v_i(x_i(\mathbf{d})) \\ &\geq \frac{1}{c} \sum_i v_i(y_i) \\ &= \frac{1}{c} SW_{OPT}(\mathbf{v}) \end{aligned}$$

as required. \square

The power of Theorem 4.3 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. An example is given in Appendix A.

4.2. Bayes-Nash Equilibria

We are now ready to bound the mixed Bayesian price of anarchy for mechanism $\mathcal{M}_1(\mathcal{A})$.

THEOREM 4.4. *Suppose \mathcal{A} is a monotone strongly loser independent allocation rule for a combinatorial allocation problem. Then the Bayesian price of anarchy of $\mathcal{M}_1(\mathcal{A})$ is at most⁸ $\frac{c}{1-e^{-c}}$ for every independent type distribution \mathbf{F} .*

We note that $\frac{c}{1-e^{-c}} \leq c(1 + \frac{2}{e^c}) = c + O(c/e^c)$. The remainder of this subsection is dedicated to the proof of Theorem 4.4.

Fix a product distribution \mathbf{F} over type profiles and let $\mathbf{b}(\cdot)$ be a (possibly mixed) Bayes-Nash equilibrium with respect to \mathbf{F} . Choose some type declaration \mathbf{v} and let $\mathbf{y}^{\mathbf{v}}$ denote an optimal allocation for \mathbf{v} . Following the proof of Theorem 4.3, we would like to bound the expected value of $\theta_i(y_i^{\mathbf{v}}, \mathbf{d}_{-i})$ with respect to $v_i(y_i^{\mathbf{v}})$ and $u_i(\mathbf{b}(\mathbf{v}))$ for each i . We encapsulate this bound in Lemma 4.6 and Corollary 4.7, below. This will allow us to use Lemma 3.3 to obtain a relation between the expected welfare of \mathcal{A} and the expected optimal welfare; this relationship is given in Lemma 4.5.

LEMMA 4.5. *Suppose that \mathcal{A} is a c -approximate monotone strongly loser independent allocation rule and that there exist constants $\gamma \geq 0$ and $\sigma_i \in [0, c]$ for $i \in [n]$ such that, whenever \mathbf{b} is a Bayes-Nash equilibrium for $\mathcal{M}_1(\mathcal{A})$, it is the case that for all i , all v_i , and all $S \subseteq M$,*

$$\mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq \gamma v_i(S) - \sigma_i \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))].$$

Then $\mathbf{E}_{\mathbf{v}}[SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})] \geq \frac{\gamma}{c} \mathbf{E}_{\mathbf{v}}[SW_{OPT}(\mathbf{v})]$.

LEMMA 4.6. *Suppose that \mathbf{b} is a Bayes-Nash equilibrium for mechanism $\mathcal{M}_1(\mathcal{A})$ and distribution \mathbf{F} . Then for all i , all v_i , and all $S \subseteq M$,*

$$\mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq v_i(S) - \left(1 + \ln \frac{v_i(S)}{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}\right) \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))].$$

Before proving Lemmas 4.5 and 4.6, let us show how they imply Theorem 4.4. We first note the following simple corollary of Lemma 4.6.

⁸In the initial conference version of this work, we presented a bound of $c + O(\log c)$ on the BPOA. Subsequently, this bound was independently improved by Lucier [Lucier 2011] to $c + O(c^2/e^c)$ and by Syrgkanis and Tardos [Syrgkanis and Tardos 2013] to $c + O(c/e^c)$. We present here a slightly modified version of the argument from Lucier [Lucier 2011], which yields the improved Syrgkanis and Tardos [Syrgkanis and Tardos 2013] bound of $c + O(c/e^c)$.

COROLLARY 4.7. *Suppose that \mathbf{b} is a Bayes-Nash equilibrium for mechanism $\mathcal{M}_1(A)$ and distribution \mathbf{F} . Then for all i , all v_i , and all $S \subseteq M$,*

$$\mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq (1 - e^{-c}) \cdot v_i(S) - c \cdot \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))].$$

PROOF. Fix agent i . By Lemma 4.6, we know

$$\mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq v_i(S) - \left(1 + \ln \frac{v_i(S)}{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}\right) \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]. \quad (3)$$

Note that if $\left(1 + \ln \frac{v_i(S)}{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}\right) \leq c$ then (3) immediately implies the desired result. We can therefore assume otherwise, and choose $\alpha > 0$ such that

$$\left(1 + \ln \frac{v_i(S)}{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}\right) = c + \alpha.$$

Rearranging, we get that $v_i(S) = e^\alpha \cdot e^{c-1} \cdot \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]$. Applying these two equalities to (3), we have

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] &\geq v_i(S) - (c + \alpha) \cdot \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))] \\ &= v_i(S) - \frac{\alpha}{e^\alpha} \cdot \frac{v_i(S)}{e^{c-1}} - c \cdot \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]. \end{aligned}$$

Since $\frac{\alpha}{e^\alpha}$ achieves its maximum value of $1/e$ at $\alpha = 1$, we can conclude that

$$\mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq v_i(S) - \frac{1}{e^c} \cdot v_i(S) - c \cdot \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]$$

as required. \square

Theorem 4.4 follows directly from Corollary 4.7 and Lemma 4.5. We next complete the proof of Theorem 4.4 by proving Lemmas 4.5 and 4.6.

Proof of Lemma 4.5 : Fix distribution \mathbf{F} over type profiles and let $\mathbf{b}(\cdot)$ be a (possibly mixed) Bayes-Nash equilibrium with respect to \mathbf{F} . Choose some type declaration \mathbf{v} and let $\mathbf{y}^{\mathbf{v}}$ denote an optimal allocation for \mathbf{v} . We know that for all $i \in [n]$ and \mathbf{v} ,

$$\mathbf{E}_{\mathbf{v}'_{-i}}[\theta_i(y_i^{\mathbf{v}}, \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \geq \gamma v_i(y_i^{\mathbf{v}}) - \sigma_i \mathbf{E}_{\mathbf{v}'_{-i}}[u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}'_{-i}))].$$

Note the distinction between \mathbf{v}'_{-i} , over which we are taking expectations, and \mathbf{v}_{-i} , which is the type profile fixed to define $y_i^{\mathbf{v}}$. Now, summing over i and taking expectation over all choices of \mathbf{v} , we have

$$\mathbf{E}_{\mathbf{v}} \left[\sum_i \mathbf{E}_{\mathbf{v}'_{-i}}[\theta_i(y_i^{\mathbf{v}}, \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \right] \geq \gamma \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(y_i^{\mathbf{v}}) \right] - \mathbf{E}_{\mathbf{v}} \left[\sum_i \sigma_i \mathbf{E}_{\mathbf{v}'_{-i}}[u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \right]. \quad (4)$$

We now consider each of the three terms in (4). First, note that

$$\mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(y_i^{\mathbf{v}}) \right] = \mathbf{E}_{\mathbf{v}}[SW_{OPT}(\mathbf{v})]. \quad (5)$$

Additionally,

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}} \left[\sum_i \sigma_i \mathbf{E}_{\mathbf{v}'_{-i}} [u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \right] &= \sum_i \sigma_i \mathbf{E}_{\mathbf{v}, \mathbf{v}'_{-i}} [u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \\
&= \mathbf{E}_{\mathbf{v}} \left[\sum_i \sigma_i u_i(\mathbf{b}(\mathbf{v})) \right] \\
&= \mathbf{E}_{\mathbf{v}} \left[\sum_i \sigma_i v_i(x_i(\mathbf{b}(\mathbf{v}))) \right] - \mathbf{E}_{\mathbf{v}, \mathbf{d}=\mathbf{b}(\mathbf{v})} \left[\sum_i \sigma_i d_i(x_i(\mathbf{b}(\mathbf{v}))) \right]
\end{aligned} \tag{6}$$

where the final equality follows from the fact that our mechanism employs a first price payment scheme. Finally,

$$\begin{aligned}
&\mathbf{E}_{\mathbf{v}} \left[\sum_i \mathbf{E}_{\mathbf{v}'_{-i}} [\theta_i(y_i^{\mathbf{v}}, \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \right] \\
&= \mathbf{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_i \theta_i(y_i^{\mathbf{v}}, \mathbf{b}_{-i}(\mathbf{v}'_{-i})) \right] \quad (\text{type independence}) \\
&\leq c \mathbf{E}_{\mathbf{v}, \mathbf{v}', \mathbf{d}'=\mathbf{b}(\mathbf{v}')} \left[\sum_i d_i'(x_i(\mathbf{d}')) \right] \quad (\text{Lemma 3.3}) \\
&= c \mathbf{E}_{\mathbf{v}, \mathbf{d}=\mathbf{b}(\mathbf{v})} \left[\sum_i d_i(x_i(\mathbf{d})) \right]
\end{aligned} \tag{7}$$

Where the final equality follows from a change of variables, since \mathbf{v} does not appear inside the expectation on the previous line. Substituting (5), (6), and (7) into (4), we conclude that

$$c \mathbf{E}_{\mathbf{v}} \left[\sum_i d_i(x_i(\mathbf{d})) \right] \geq \gamma \mathbf{E}_{\mathbf{v}} [SW_{OPT}(\mathbf{v})] - \mathbf{E}_{\mathbf{v}} \left[\sum_i \sigma_i v_i(x_i(\mathbf{b}(\mathbf{v}))) \right] + \mathbf{E}_{\mathbf{v}, \mathbf{d}=\mathbf{b}(\mathbf{v})} \left[\sum_i \sigma_i d_i(x_i(\mathbf{b}(\mathbf{v}))) \right]$$

and hence

$$\begin{aligned}
\gamma \mathbf{E}_{\mathbf{v}} [SW_{OPT}(\mathbf{v})] &\leq \mathbf{E}_{\mathbf{v}} \left[\sum_i \sigma_i v_i(x_i(\mathbf{b}(\mathbf{v}))) \right] + \mathbf{E}_{\mathbf{v}, \mathbf{d}=\mathbf{b}(\mathbf{v})} \left[\sum_i (c - \sigma_i) d_i(x_i(\mathbf{d})) \right] \\
&\leq \mathbf{E}_{\mathbf{v}} \left[\sum_i \sigma_i v_i(x_i(\mathbf{b}(\mathbf{v}))) \right] + \mathbf{E}_{\mathbf{v}} \left[\sum_i (c - \sigma_i) v_i(x_i(\mathbf{b}(\mathbf{v}))) \right] \\
&= \mathbf{E}_{\mathbf{v}} \left[\sum_i c v_i(x_i(\mathbf{b}(\mathbf{v}))) \right] \\
&= c \mathbf{E}_{\mathbf{v}} [SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})]
\end{aligned}$$

where in the second inequality we used Corollary 4.2 plus the fact that $(c - \sigma_i) \geq 0$ for all i . Rearranging yields

$$\mathbf{E}_{\mathbf{v}} [SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})] \geq \frac{\gamma}{c} \mathbf{E}_{\mathbf{v}} [SW_{OPT}(\mathbf{v})]$$

as required. \square

Proof of Lemma 4.6 : Fix any i , v_i , and S . Since $\theta_i(S, \mathbf{d}_{-i}) \geq 0$ for all \mathbf{d}_{-i} , we have that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] &\geq \int_0^{v_i(S)} \Pr[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i})) > z] dz \\ &= v_i(S) - \int_0^{v_i(S)} \Pr[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \leq z] dz. \end{aligned}$$

Recall that $b_i(v_i)$ must maximize the expected utility of agent i . Choose any $z \geq 0$ and consider the alternative strategy d_i which places a single-minded bid of z on set S . Then since $b_i(v_i)$ is an optimal strategy, we have that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))] &\geq \mathbf{E}_{\mathbf{v}_{-i}}[u_i(d_i, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \\ &= (v_i(S) - z) \Pr[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \leq z] \end{aligned}$$

where the equality follows since any single minded bid above the critical value for S insures that S will be won, as a consequence of monotonicity. We conclude that

$$\Pr[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \leq z] \leq \frac{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}{(v_i(S) - z)}$$

for all $0 \leq z < v_i(S)$. We also know that $\Pr[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \leq z] \leq 1$ for all z . Write $r = v_i(S) - \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]$. We then conclude that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] &\geq v_i(S) - \int_0^r \frac{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}{(v_i(S) - z)} dz - \int_r^{v_i(S)} 1 dz \\ &= v_i(S) - \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))] \int_{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}^{v_i(S)} \frac{1}{y} dy - \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))] \\ &= v_i(S) - \left(1 + \ln \frac{v_i(S)}{\mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))]}\right) \mathbf{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{b}(\mathbf{v}))] \end{aligned}$$

as required. \square

COROLLARY 4.8. (of proof) *The same bound on price of anarchy applies to coarse correlated equilibrium.*

PROOF. That such price of anarchy bounds can be applied to coarse correlated equilibria in the full information setting was initially observed by Roughgarden [Roughgarden 2009]. Specifically, in the proof of Theorem 4.4, all occurrences of $\mathbf{E}_{\mathbf{v}, \mathbf{d}=\mathbf{b}(\mathbf{v})}$ can be replaced by $\mathbf{E}_{\mathbf{d} \sim (d_i', \omega_{-i})}$, resulting in a bound on the coarse correlated PoA. \square

It may be tempting to conjecture that the (exponentially small) loss in approximation factor in Theorem 4.4 is simply an artifact of the analysis, and that the Bayesian price of anarchy of $\mathcal{M}_1(\mathcal{A})$ is actually c . However, we now show by way of an example that this loss is necessary; that is, there exist instances in which the mixed price of anarchy (and hence the Bayesian price of anarchy) is strictly greater than c .

PROPOSITION 4.9. *For any $c \geq 2$, there is a combinatorial allocation problem \mathcal{P} and a non-adaptive greedy algorithm A such that A is a c -approximation for \mathcal{P} , and the mixed price of anarchy for $\mathcal{M}_1(\mathcal{A})$ is at least $c + c^2/e^{4c} = c + \Omega(c^2/e^{4c})$.*

PROOF. We begin by describing our combinatorial allocation problem. We choose a parameter $k > c$ that will be fixed later. Our auction has $ck + k$ objects, which we label a_{ij} for $i \in [k], j \in [c]$ and b_i for $i \in [k]$. There are $4k$ agents, labelled A_i, B_i, C_i , and D_i for $i \in [k]$. Our feasibility constraints are as follows. Each agent B_i or C_i can receive only set $\{a_{i1}\}$ or \emptyset . Each agent D_i can receive set $\{a_{i1}, a_{k1}\}$ or \emptyset . Each agent A_i can receive either set $\{a_{i1}, a_{i2}, \dots, a_{ic}\}$, set $\{b_i\}$, or \emptyset . Under these restrictions, an allocation is feasible if each object is assigned to at most one agent.

Let \mathcal{A} be the non-adaptive greedy algorithm that orders bids by density: i.e. with priority function $r(i, S, v) = v/|S|$ when S is a feasible set for agent i . We claim that when $c \geq 2$, this algorithm obtains a c -approximation for the above combinatorial auction. To see this, note that the (unique) set that can be allocated to any agent B_i, C_i , or D_i intersects sets of size at most c times larger, so if the greedy algorithm allocates to one of these agents for a value of v , the total value of intersecting sets in the optimal solution is at most cv . On the other hand, if the greedy algorithm allocates $\{b_i\}$ to agent A_i , this conflicts only with the allocation of set $\{a_{i1}, \dots, a_{ic}\}$ to agent A_i , which again has value at most c times greater. Finally, suppose that the greedy algorithm allocates set $\{a_{i1}, \dots, a_{ic}\}$ to agent A_i , say with value vc (i.e. value density v). This allocation can conflict only with a single allocation to an agent B_i, C_i , or D_i plus an allocation of $\{b_i\}$ to agent A_i , which comprises a total of at most 3 objects. Since the greedy algorithm allocates by density, the total value of the conflicted bids is at most $3v$. Since $c \geq 2$, we conclude that the allocation of $\{a_{i1}, \dots, a_{ic}\}$ to agent A_i is within a factor of c of the value of any intersecting sets in the optimal allocation.

Consider now the following instance of this problem, specified by the following agent types.

- For $1 \leq i \leq k - 1$, agent A_i desires $\{a_{i1}, a_{i2}, \dots, a_{ic}\}$ for value $k - i$ and $\{b_i\}$ for value 0.
- Agent A_k desires $\{a_{k1}, a_{k2}, \dots, a_{kc}\}$ for value k and $\{b_k\}$ for value 1.
- For $1 \leq i \leq k$, agents B_i and C_i both desire set $\{a_{i1}\}$ for value $(k - i)/c$.
- For $1 \leq i \leq k$, agent D_i desires set $\{a_{i1}, a_{k1}\}$ for value $2(k - i)/c$.

Note that agent A_k has a value density of k/c for the desired set $\{a_{k1}, \dots, a_{kc}\}$, and each agent A_i with $i < k$ has value density $(k - i)/c$ for desired set $\{a_{i1}, \dots, a_{ic}\}$. Also, the agents B_i, C_i , and D_i have a value density of $(k - i)/c$ for their desired sets.

We will suppose that \mathcal{A} applies the following fixed tie-breaking rules. For any i , \mathcal{A} will break a tie between agents A_i, B_i, C_i , and/or D_i first in favour of D_i , then in favour of B_i , then A_i . We can also assume that \mathcal{A} breaks ties between multiple desired sets for agent A_i in favour of $\{b_i\}$. Finally, \mathcal{A} will favour allocating non-empty sets over allocating the empty set (e.g., if an agent declares the zero valuations).

We now describe a mixed Nash equilibrium for this problem instance. Each agent A_i declares the zero valuation. Each agent B_i and C_i declares his valuation truthfully. Each agent D_i will declare his valuation truthfully with some probability p_i , and will otherwise declare the zero valuation. We choose $p_i = \frac{1}{i+1}$.

What is the outcome when agents bid in this way? First, each agent A_i is allocated set $\{b_i\}$ (due to our assumed tie-breaking). For the items a_{ij} , only items with $j = 1$ will be allocated. For $i < k$, if agents D_1, \dots, D_{i-1} declare the zero allocation and D_i does not, then object a_{1i} will be allocated to D_i . If not, then item a_{1i} will be allocated to agent B_i . Item a_{k1} will be allocated to D_i where i is the smallest such that D_i does not declare the zero valuation, or B_k if D_1, \dots, D_k all declare the zero valuation.

We now argue that this distribution of declarations is indeed a mixed Nash equilibrium. With probability 1, no agent B_i, C_i , or D_i can obtain positive utility from any declaration (since their desired sets conflict with other bids of the same value density),

so their distributions over declarations that obtain utility 0 are necessarily optimal. Furthermore, for each $i < k$, agent A_i cannot obtain positive utility so his bidding strategy is also optimal. Agent A_k obtains utility 1; his only hope for obtaining more utility is to declare a value less than $k - 1$ for set $\{a_{k1}, \dots, a_{kc}\}$. However, if he declares some value $k - z$ with $z > 1$, say with $x = \lceil z \rceil$, then he can win his desired set only if bidders D_1, \dots, D_{x-1} all bid the zero valuation, since otherwise an agent D_j with $j < x$ would win his desired set, blocking the bid by agent A_k . The probability that bidders D_1, \dots, D_{x-1} all declare the zero valuation is $\frac{1}{2} \cdot \frac{2}{3} \cdots \frac{x-1}{x} = \frac{1}{x} \leq \frac{1}{z}$. Thus, for any z , agent A_k can obtain utility z with probability at most $1/z$, for an expected utility of at most 1. The given declaration by agent A_k is therefore optimal.

We will now bound the social efficiency of this equilibrium. The optimal obtainable welfare is $k + \sum_{i=1}^{k-1} (k - i) = \frac{1}{2}k(k + 1)$, by allocating set $\{a_{i1}, \dots, a_{ic}\}$ to agent A_i for all i . In the equilibrium we've described, object b_k is allocated to agent A_k for a value of 1 and each object a_{i1} for $i < k$ is allocated to either B_i or D_i at a per-item value of $(k - i)/c$. For each $i < k$, object a_{i1} will be allocated to bidder D_i precisely if bidders D_1, \dots, D_{j-1} declare the zero valuation but D_i does not, which occurs with probability $\frac{1}{i(i+1)}$. Object a_{i1} will be allocated to either B_k or D_k with the remaining probability, which is $\frac{1}{k}$. Noting that each of B_i and D_i has a per-item value of $(k - i)/c$ for their desired sets, we conclude that the expected total value obtained is

$$\begin{aligned} & 1 + \sum_{i < k} \frac{k - i}{c} + \sum_{i < k} \frac{1}{i(i+1)} \cdot \frac{k - i}{c} + \frac{1}{k} \cdot \frac{k - k}{c} \\ &= 1 + \frac{1}{c} \left[\frac{1}{2}(k^2 - k) + k - \sum_{i < k} \frac{1}{i+1} - 1 \right] \\ &= 1 + \frac{1}{c} \left[\frac{1}{2}(k^2 + k) - H_k \right] \end{aligned}$$

where H_k is the k th harmonic number.

We conclude that the mixed price of anarchy for this mechanism is at least

$$\frac{\frac{1}{2}(k^2 + k)}{1 + \frac{1}{c} \left[\frac{1}{2}(k^2 + k) - H_k \right]} > c \left(\frac{k^2 + k}{k^2 + k + 2c - 2 \ln k} \right)$$

where we used the fact that $H_k > \ln k$. Choose $k = e^{2c}$. Then our mechanism has mixed price of anarchy at least

$$c \left(\frac{e^{4c} + e^{2c}}{e^{4c} + e^{2c} - 2c} \right) > c \left(\frac{e^{4c}}{e^{4c} - c} \right) > c \left(1 + \frac{c}{e^{4c}} \right)$$

as required. \square

4.3. Correlated Types

Recall that our bound for Bayesian Price of Anarchy required that agent types be distributed independently. We now provide an alternative (weaker) bound that holds even if agent types are arbitrarily correlated. The key to the new analysis is in considering a deviating behaviour for each agent that does not depend on the other agents' types. The particular deviation we will consider is that of bidding half of one's true value for every set. Our analysis will additionally require that the underlying allocation algorithm is a fixed order greedy algorithm.

THEOREM 4.10. *Suppose \mathcal{A} is a c -approximate non-adaptive greedy algorithm for a combinatorial allocation problem. Then $\mathcal{M}_1(\mathcal{A})$ has Correlated Bayesian Price of Anarchy at most $4c$, for any type distribution \mathbf{F} .*

The key to this result lies in the following lemma.

LEMMA 4.11. *Suppose \mathcal{A} is a c -approximate non-adaptive greedy algorithm for a combinatorial allocation problem. Then for all type profiles \mathbf{v} and all strategy profiles $\mathbf{b}(\cdot)$,*

$$\sum_i u_i(v_i/2, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \geq \frac{1}{2c} SW_{OPT}(\mathbf{v}) - SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v}).$$

PROOF. Let \mathbf{y} denote the optimal allocation for type profile \mathbf{v} . Choose agent i , and consider the outcome of \mathcal{A} on input profile $(v_i/2, \mathbf{b}_{-i}(\mathbf{v}_{-i}))$. Let $x_i = \mathcal{A}_i(v_i/2, \mathbf{b}_{-i}(\mathbf{v}_{-i}))$. Note that it must either be that $\theta_i(y_i, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \geq \frac{1}{2}v_i(y_i)$ or not. In the latter case, agent i must obtain some allocation x_i with $r(i, x_i, v_i(x_i)/2) \geq r(i, y_i, v_i(y_i)/2)$. Since \mathcal{A} is a non-adaptive greedy algorithm, this then implies that $v_i(x_i) \geq \frac{1}{c}v_i(y_i)$, since otherwise \mathcal{A} would obtain less than a $\frac{1}{c}$ fraction of the optimal social welfare on the input in which agent i places bids only on sets x_i and y_i , and all other agents bid 0.

We conclude that for all i , either $\theta_i(y_i, \mathbf{b}_{-i}(\mathbf{v}_{-i})) > \frac{1}{2}v_i(y_i)$ or else $v_i(x_i) \geq \frac{1}{c}v_i(y_i)$. Let $N = \{i \mid \theta_i(y_i, \mathbf{b}_{-i}(\mathbf{v}_{-i})) > \frac{1}{2}v_i(y_i)\}$ be the set of agents for which the former condition holds. We then note that

$$\sum_{i \in N} \frac{1}{2}v_i(y_i) < \sum_{i \in N} \theta_i(y_i, \mathbf{b}_{-i}(\mathbf{v}_{-i})) \leq cSW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{b}(\mathbf{v})) \leq cSW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})$$

where the second inequality is due to Lemma 3.3 and the third is due to Lemma 4.1. Furthermore, since $v_i(x_i) \geq \frac{1}{c}v_i(y_i)$ for all $i \notin N$, we have

$$\sum_{i \notin N} \frac{1}{2}v_i(y_i) \leq \sum_{i \notin N} \frac{c}{2}v_i(x_i(v_i/2, \mathbf{b}_{-i}(\mathbf{v}_{-i}))) \leq c \sum_i u_i(v_i/2, \mathbf{b}_{-i}(\mathbf{v}_{-i}))$$

where the second inequality follows because we are using the first-price payment scheme. Combining these inequalities yields

$$\sum_i u_i(v_i/2, \mathbf{b}_{-i}(\mathbf{v}_{-i})) + SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v}) \geq \frac{1}{2c} SW_{OPT}(\mathbf{v})$$

as required. \square

Theorem 4.10 now follows easily from Lemma 4.11. Recall that Lemma 4.11 holds for all strategy profiles, not just strategies in equilibrium. If we take \mathbf{b} to be an equilibrium

profile under type distribution \mathbf{F} , then

$$\begin{aligned}
E_{\mathbf{v}}[SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})] &\geq E_{\mathbf{v}} \left[\sum_i u_i(\mathbf{b}(\mathbf{v})) \right] \\
&= \sum_i E_{v_i} E_{\mathbf{v}_{-i}|v_i} [u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \\
&\geq \sum_i E_{v_i} E_{\mathbf{v}_{-i}|v_i} \left[u_i \left(\frac{v_i}{2}, \mathbf{b}_{-i}(\mathbf{v}_{-i}) \right) \right] \\
&= E_{\mathbf{v}} \left[\sum_i u_i \left(\frac{v_i}{2}, \mathbf{b}_{-i}(\mathbf{v}_{-i}) \right) \right] \\
&\geq E_{\mathbf{v}} \left[\frac{1}{2c} \sum_i SW_{OPT}(\mathbf{v}) - SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v}) \right] \quad (\text{Lemma 4.11})
\end{aligned}$$

from which we conclude that

$$E_{\mathbf{v}}[SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})] \geq \frac{1}{4c} E_{\mathbf{v}} \left[\sum_i OPT(\mathbf{v}) \right]$$

completing the proof of Theorem 4.10.

5. CRITICAL-PRICE MECHANISMS

We begin by studying the performance of critical (i.e. second) price mechanisms at equilibrium. The mechanism we study is $\mathcal{M}_2(\mathcal{A})$, which is defined with respect to an arbitrary monotone strongly loser-independent algorithm \mathcal{A} . Recall that $\mathcal{M}_2(\mathcal{A})$ proceeds by first collecting a declaration profile from the agents, then passing the observed declarations to \mathcal{A} as input. The mechanism returns the allocation provided by \mathcal{A} as output, and charges each agent his critical value for the set received (computed via additional calls to \mathcal{A} ; see Section 5.4).

We will show that every Bayes-Nash equilibria of $\mathcal{M}_2(\mathcal{A})$ has a social welfare guarantee nearly matching that of the original algorithm \mathcal{A} . This result requires that we make an assumption on the bidding strategies applied by the users; namely, that they do not *overbid*, meaning that they do not bid more than their true value on any given set S . This overbidding assumption is necessary to exclude certain degenerate equilibria, such as one agent making an infinitely large bid on the set of all objects and other bidders bidding 0. We note that such assumptions are reasonable in general; even the truthful Vickrey auction of a single item requires a no-overbidding assumption to bound the efficiency of the outcome at equilibrium. In Section 5.3 we discuss ways to relax this assumption by modifying the mechanism slightly.

5.1. Bayes-Nash Equilibria

We begin by analyzing the Bayesian price of anarchy for the critical price mechanism $\mathcal{M}_2(\mathcal{A})$. Given that agents will not overbid, a simple modification of Theorem 4.4 yields a result for BNE under critical prices.

THEOREM 5.1. *Suppose \mathcal{A} is a c -approximate monotone strongly loser independent allocation rule, and that \mathbf{b} is a Bayes-Nash equilibrium of $\mathcal{M}_2(\mathcal{A})$ in which agents do not overbid. Then the expected welfare when agents declare according to \mathbf{b} is a $(c + 1)$ -approximation to the expected optimal welfare.*

LEMMA 5.2. *Suppose that \mathbf{b} is a Bayes-Nash equilibrium for mechanism $\mathcal{M}_2(\mathcal{A})$ and distribution \mathbf{F} . Then for all i , all v_i , and all $S \subseteq M$,*

$$\mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq v_i(S) - \mathbf{E}_{\mathbf{v}_{-i}}[v_i(x_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})))]$$

PROOF. Choose any i , v_i , and S . Let d_i be a single-minded declaration for set S at value $v_i(S)$, and consider a strategy under which agent i declares d_i when his type is v_i . Under this strategy, the expected utility of agent i with type v_i is

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}}[u_i(d_i, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] &\geq \mathbf{E}_{\mathbf{v}_{-i}}[\max\{v_i(S) - \theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i})), 0\}] \\ &\geq v_i(S) - \mathbf{E}_{\mathbf{v}_{-i}}[\theta_i(S, \mathbf{b}_{-i}(\mathbf{v}_{-i}))]. \end{aligned} \quad (8)$$

Since b_i is an equilibrium strategy for agent i , it must be that

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}}[u_i(d_i, \mathbf{b}_{-i}(\mathbf{v}_{-i}))] &\leq \mathbf{E}_{\mathbf{v}_{-i}}[u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \\ &\leq \mathbf{E}_{\mathbf{v}_{-i}}[v_i(x_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i})))]. \end{aligned} \quad (9)$$

Combining equations (8) and (9) leads to the desired result. \square

Following the proof of Theorem 4.4, we conclude that for all equilibria \mathbf{b} , if we write $\mathbf{y}^{\mathbf{v}}$ for an optimal allocation for any given type profile \mathbf{v} , then

$$\begin{aligned} \mathbf{E}_{\mathbf{v}} \left[\sum_i \mathbf{E}_{\mathbf{v}'_{-i}}[\theta_i(y_i^{\mathbf{v}}, \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \right] &\geq \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(y_i^{\mathbf{v}}) \right] \\ &\quad - \mathbf{E}_{\mathbf{v}} \left[\sum_i \mathbf{E}_{\mathbf{v}'_{-i}}[v_i(x_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}'_{-i})))] \right]. \end{aligned} \quad (10)$$

Just as in the proof of Theorem 4.4, we obtain the bounds

$$\begin{aligned} \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i(y_i^{\mathbf{v}}) \right] &= \mathbf{E}_{\mathbf{v}}[SW_{OPT}(\mathbf{v})], \\ \mathbf{E}_{\mathbf{v}} \left[\sum_i \mathbf{E}_{\mathbf{v}'_{-i}}[v_i(x_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}'_{-i})))] \right] &= \mathbf{E}_{\mathbf{v}}[SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})], \\ \mathbf{E}_{\mathbf{v}} \left[\sum_i \mathbf{E}_{\mathbf{v}'_{-i}}[\theta_i(y_i^{\mathbf{v}}, \mathbf{b}_{-i}(\mathbf{v}'_{-i}))] \right] &\leq c \mathbf{E}_{\mathbf{v}}[SW(\mathcal{A}(\mathbf{b}(\mathbf{v})), \mathbf{v})] \end{aligned}$$

which, taken together with (10), completes the proof of Theorem 5.1. Note that when deriving the last inequality above, we do not invoke Lemma 4.1 (as in the proof of Theorem 4.4); instead, we use the assumption that agents do not overbid. \square

In precisely the same way as for the first-price mechanism, the bound on the price of anarchy also extends to coarse correlated equilibrium.

COROLLARY 5.3. *(of proof) The bound of $(c + 1)$ on the price of anarchy applies also to coarse correlated equilibrium.*

We next show that this gap between the approximation factor of the original algorithm and the price of anarchy of the critical price mechanism is required for large c . For any $c \geq 1$ we exhibit a combinatorial allocation problem and a non-adaptive greedy algorithm \mathcal{A} such that the approximation factor of \mathcal{A} is $c + \frac{1}{c}$ but the (pure) price of anarchy of $\mathcal{M}_2(\mathcal{A})$ is $c + 1$. This leads us to conclude that, in general, the bound in Theorem 5.1 cannot be improved beyond $c + 1 - \theta(\frac{1}{c})$.

PROPOSITION 5.4. *For any $c \geq 1$, there is a combinatorial allocation problem \mathcal{P} and a non-adaptive greedy algorithm \mathcal{A} such that \mathcal{A} is a $(c + \frac{1}{c})$ -approximation for \mathcal{P} , and the pure price of anarchy for $\mathcal{M}_2(\mathcal{A})$ is $c + 1$.*

PROOF. Consider a combinatorial auction problem with two objects a, b and two players, under the restriction that each player can be allocated at most one object and player 2 cannot be allocated object b . Algorithm \mathcal{A} will be the following non-adaptive greedy algorithm: if $v_1(a) \geq \frac{1}{c}v_2(a)$ and $v_1(a) \geq cv_1(b)$ then allocate a to player 1 and \emptyset to player 2; otherwise allocate b to player 1 and a to player 2. Note that this is a $c + \frac{1}{c}$ approximation algorithm, since whenever the algorithm allocates a to player 1 we have $v_2(a) + v_1(b) \leq (c + \frac{1}{c})v_1(a)$, and whenever the algorithm allocates a to player 2 we have $v_1(a) \leq c(v_1(b) + v_2(b))$.

Consider the mechanism $\mathcal{M}_2(\mathcal{A})$, and suppose that the agents have a type profile in which $v_1(a) = v_1(b) = 1$ and $v_2(a) = c$. Then the declaration profile $d_1(a) = 1, d_1(b) = 0$, and $d_2(a) = 0$ is in equilibrium, since agent 1 cannot improve upon his utility of 1 and agent 2 cannot affect the outcome without paying at least $\theta_2(a, d_1) = c$, for a utility of 0. The social welfare at this equilibrium is 1, but a total of $c + 1$ is possible by allocating a to player 2 and allocating b to player 1. Thus the price of anarchy for $\mathcal{M}_2(\mathcal{A})$ is at least $c + 1$. \square

5.2. Correlated Types

Theorem 5.1 requires that agent types be distributed independently. As with the first-price mechanism, we can provide a somewhat weaker bound that holds even when agent types are arbitrarily correlated. And as in Theorem 5.1 this result additionally requires that the underlying allocation algorithm is a non-adaptive greedy algorithm.

THEOREM 5.5. *Suppose \mathcal{A} is a c -approximate non-adaptive greedy algorithm for a combinatorial allocation problem, and that agents do not overbid. Then $\mathcal{M}_2(\mathcal{A})$ has Correlated Bayesian Price of Anarchy at most $4c$, for any type distribution \mathbf{F} .*

The proof of Theorem 5.5 follows that of Theorem 4.10 almost exactly. The sole difference is that the invocation of Lemma 4.1 in the proof of Theorem 4.10 is replaced by an appeal to the no-overbidding assumption. We omit the details for brevity.

5.3. Overbidding and Restricted Expressiveness

Our analysis to this point made use of a no-overbidding assumption, which states that no agent will place a bid larger than her true value on any given set. However, our use of the no-overbidding assumption is marred by the fact that a restriction to no-overbidding strategies is not always rational when agents have complete confidence about their opponents' type distributions. As the following example shows, an agent may be *strictly* better off by overbidding, even in a full-information setting. In other words, a strategy with overbidding is not necessarily dominated.

Example 5.1. 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 greedy algorithm that orders bids by value. 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 (as any optimal randomized strategy has only optimal strategies in its support). Suppose agent 1 does not overbid and declares at most 2 for object b . If he also

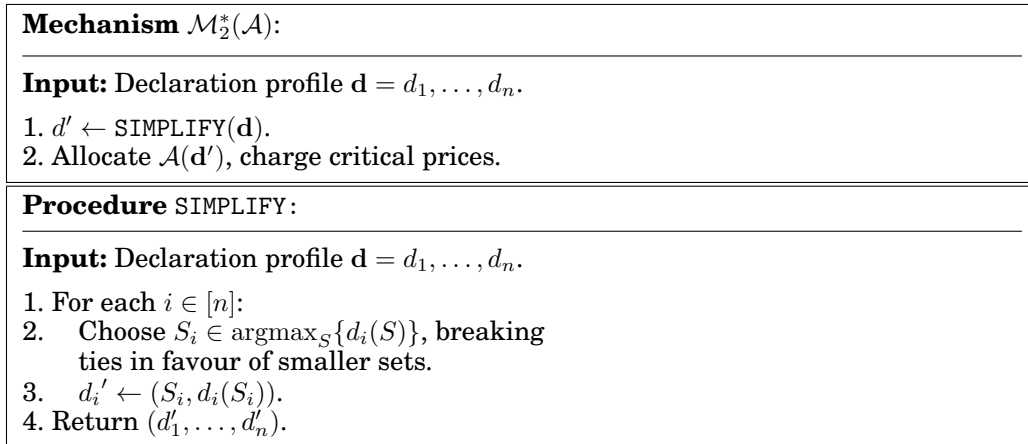


Fig. 2. Simplifying declarations in a critical price mechanism.

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 now show that if we modify mechanism $\mathcal{M}_2(\mathcal{A})$ by effectively limiting the expressiveness of the bids made by the agents, then we obtain the same efficiency bounds at equilibria but furthermore guarantee that any bidding strategy that involves overbidding is dominated. Thus, as long as agents avoid dominated strategies (a very mild assumption), all equilibria of rational play lead to approximately efficient outcomes.

For a monotone strongly loser independent allocation rule \mathcal{A} , the modified mechanism $\mathcal{M}_2^*(\mathcal{A})$ is as described in Figure 2. Mechanism $\mathcal{M}_2^*(\mathcal{A})$ proceeds by first simplifying the declaration given by each agent, then passing the simplified declarations to algorithm \mathcal{A} . The resulting allocation is paired with a payment scheme that charges critical prices.

The simplification process SIMPLIFY essentially converts any declaration into a single-minded declaration (and does not affect declarations that are already single-minded). We can therefore assume without loss of generality that agents always make single-minded declarations to this mechanism, as additional information is not used.⁹

Fix a particular combinatorial auction problem and type profile \mathbf{v} , and let \mathcal{A} be an arbitrary strongly loser-independent approximation algorithm. Since \mathbf{v} is fixed, we can think of a strategy for each agent i as a declaration $d_i \in V_i$. Let \mathbf{d} be a declaration profile; we suppose each d_i is a single-minded bid for set S_i (and, in general, we will write S_i for the desired set in declaration d_i). We draw the following conclusion about the bidding choices of rational agents.

LEMMA 5.6. *Let \mathcal{A} be monotone strongly loser independent allocation rule, and fix type profile \mathbf{v} . Then for each agent i , a single-minded declaration d_i for set S_i is an undominated strategy for mechanism $\mathcal{M}_2^*(\mathcal{A})$ if and only if $d_i(S_i) = v_i(S_i)$.*

⁹We note, however, that this is not the same as assuming that agents are single-minded; our results hold for bidders with general private valuations.

PROOF. Fix some \mathbf{d}_{-i} and suppose d_i is a single-minded declaration for set S_i . On input (d_i, \mathbf{d}_{-i}) , mechanism $\mathcal{M}_2^*(\mathcal{A})$ either allocates S_i or \emptyset to agent i . Thus agent i 's utility for declaring d_i , $u_i(d_i, \mathbf{d}_{-i})$, is $v_i(S_i) - \theta_i(S_i, \mathbf{d}_{-i})$ when $d_i(S_i) > \theta_i(S_i, \mathbf{d}_{-i})$, and 0 otherwise (where θ_i denotes critical prices with respect to $\mathcal{M}_2^*(\mathcal{A})$). A declaration of $d_i(S_i) = v_i(S_i)$ therefore maximizes $u_i(d_i, \mathbf{d}_{-i})$ for all \mathbf{d}_{-i} .

Next suppose that $d_i(S_i) \neq v_i(S_i)$; we will show that d_i is dominated. Let d_i' be the single-minded declaration for S_i at value $v_i(S_i)$. Suppose there is some \mathbf{d}_{-i} such that $\theta_i^{\mathcal{A}}(S_i, \mathbf{d}_{-i})$ lies strictly between $d_i(S_i)$ and $v_i(S_i)$. For simplicity we will assume such a \mathbf{d}_{-i} exists; handling the general case requires only a technical extension of notation¹⁰. Then if $d_i(S_i) < v_i(S_i)$, then $u_i(d_i', \mathbf{d}_{-i}) > 0 = u_i(d_i, \mathbf{d}_{-i})$. Otherwise, if $d_i(S_i) > v_i(S_i)$, then $u_i(d_i', \mathbf{d}_{-i}) \geq 0 > u_i(d_i, \mathbf{d}_{-i})$. Thus, in either case, we have $u_i(d_i', \mathbf{d}_{-i}) > u_i(d_i, \mathbf{d}_{-i})$, and therefore declaration d_i' strictly dominates declaration d_i . \square

Given Lemma 5.6, we can analyze the efficiency of equilibria of $\mathcal{M}_2^*(\mathcal{A})$ in a manner identical to $\mathcal{M}_2(\mathcal{A})$. Rather than explicitly assuming that agents do not overbid, Lemma 5.6 implies that they will not.

THEOREM 5.7. *Suppose \mathcal{A} is a c -approximate monotone strongly loser independent allocation rule, and that $\mathbf{b}(\cdot)$ is a Bayes-Nash equilibrium of $\mathcal{M}_2^*(\mathcal{A})$. Then the expected welfare when agents declare according to \mathbf{b} is a $(c + 1)$ approximation to the expected optimal welfare.*

5.4. Calculating Critical Prices

For many allocation algorithms (such as all of the algorithms discussed in Section 2.5), 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 by multiples of ϵ , critical prices can be determined efficiently. If valuation space is continuous, then our interpretation is that any equilibrium for the (exact) critical-price mechanism will be an (additive) ϵ -approximate equilibrium for a mechanism that uses ϵ -approximate critical prices.

We now describe the procedure for determining critical prices in more detail. Fix greedy allocation rule \mathcal{A} , agent i , and declarations \mathbf{d} . Suppose that $\mathcal{A}_i(d_i, \mathbf{d}_{-i}) = S$. We wish to resolve the value of $\theta_i(S, \mathbf{d}_{-i})$ in the range $[0, d_i(S)]$ using binary search in the following way. For all $z \geq 0$, write d_i^z for the single-minded declaration for set S at value z . Given query value $z \in [0, d_i(S)]$, we check if $\mathcal{A}_i(d_i^z, \mathbf{d}_{-i}) = S$. If so, decrease the value of z ; otherwise, increase the value of z . Since \mathcal{A} is monotone, we have that $\mathcal{A}_i(d_i^z, \mathbf{d}_{-i}) = S$ if and only if $z > \theta_i(S, \mathbf{d}_{-i})$. This procedure resolves the value of v to within ϵ in $O(\log d_i(S)/\epsilon)$ iterations. Thus, for any given input to mechanism $\mathcal{M}_2(\mathcal{A})$, the critical prices for all agents' allocated sets can be found in $O(n \log(v_{max}/\epsilon))$ invocations of algorithm \mathcal{A} , where $v_{max} = \max_{i,S} d_i(S)$.

¹⁰If $\theta_i^{\mathcal{A}}(S_i, \mathbf{d}_{-i})$ never lies between $d_i(S_i)$ and $v_i(S_i)$ for any \mathbf{d}_{-i} , then $\mathcal{M}_{\mathcal{A}}(d_i, \mathbf{d}_{-i}) = \mathcal{M}_{\mathcal{A}}(d_i', \mathbf{d}_{-i})$ for all \mathbf{d}_{-i} , so d_i and d_i' are equivalent strategies. We can therefore think of d_i as being "the same" as a single-minded declaration for S_i at value $v_i(S_i)$. We will ignore this technical issue for the remainder of the proof, in the interest of clarity.

6. REPEATED AUCTIONS AND REGRET MINIMIZATION

Up to this point, we have considered the performance of mechanisms for one-shot allocation problems. We now turn to repeated auctions. In this section, we focus on agents that apply regret minimizing strategies. We consider an instance of a combinatorial allocation problem that proceeds in rounds. The problem will be resolved by a direct revelation mechanism \mathcal{M} , say with allocation algorithm \mathcal{A} , which independently executes on each round of the auction. As before, we will tend to write \mathbf{x} for the allocation rule associated with algorithm \mathcal{A} .

We assume that neither the agent types nor the mechanism changes between rounds of the auction. When the agents have types \mathbf{v} and $D = (\mathbf{d}^1, \mathbf{d}^2, \dots, \mathbf{d}^T, \dots)$ is a sequence of declared valuation profiles, we let D_T denote the length- T prefix of D and we write $SW_{\mathcal{A}}(D_T) = \frac{1}{T} \sum_t SW(\mathbf{x}(\mathbf{d}^t), \mathbf{v})$ for the average welfare obtained over the T declarations in D_T . We will sometimes replace subscript \mathcal{A} by \mathcal{M} , in which case the social welfare is for the allocation rule of \mathcal{M} .

Declaration sequence $D = (\mathbf{d}^0, \mathbf{d}^1, \dots)$ *minimizes external regret* for agent i if, for any fixed declaration d_i , the sequence of finite prefixes satisfies $\sum_{t=1}^T u_i(d_i^t, \mathbf{d}_{-i}^t) \geq \sum_t u_i(d_i, \mathbf{d}_{-i}^t) + o(T)$. That is, as T increases, the per-round utility of agent i approaches the utility of the optimal fixed strategy in hindsight. The *price of total anarchy* [Blum et al. 2008] is the worst-case ratio between optimal welfare and the welfare of a declaration that minimizes external regret; that is,

$$\sup_{\mathbf{v}, D} \frac{SW_{opt}(\mathbf{v})}{SW(\mathcal{M}(D), \mathbf{v})}.$$

THEOREM 6.1. *For any c -approximate monotone strongly loser independent allocation rule \mathcal{A} , mechanism $\mathcal{M}_1(\mathcal{A})$ (respectively, $\mathcal{M}_2^*(\mathcal{A})$) has price of total anarchy at most $\frac{c}{1-e^{-c}}$ (resp. $c+1$).*

PROOF. As observed by Blum and Mansour [Blum and Mansour 2007] and Roughgarden [Roughgarden 2009; 2012], in the full information setting, price of total anarchy is equal to price of anarchy with respect to coarse correlated equilibrium. The result then follows immediately from Corollaries 4.8 and 5.3. \square

One can suppose that each bidder employs an algorithm to determine which declaration to make at time t , given the bidding history up to time t . We say that such an algorithm minimizes regret if employing the algorithm results in a sequence that minimizes external regret for the employing agent, for any bidding behavior of the other agents. Our interest in external regret minimization is motivated by the existence of simple and efficient bidding algorithms for minimizing regret. Indeed, the price of total anarchy captures mechanism performance when agents apply reasonable learning techniques over the course of repeated participation. In this sense we feel that this analysis is predictive of outcomes that would be observed in practice in a repeated auction. The remainder of this section will be dedicated to elaborating on this point of computational tractability.

The standard algorithmic approach to minimizing external regret is the “follow the perturbed leader” (FPL) algorithm [Hannan 1957; Kalai and Vempala 2005], which requires time and space polynomial in the number of actions that can be taken by an agent. Note that, in general, an action in a combinatorial auction mechanism corresponds to a declared valuation, of which there are super-polynomially many. So FPL is not immediately applicable as an approach for arbitrary mechanisms.

We note, however, that the mechanism $\mathcal{M}_2^*(\mathcal{A})$ (for a monotone strongly loser independent allocation rule \mathcal{A}) has some desirable properties that make it well-suited to the application of FPL techniques. Recall that,

for this mechanism, it suffices to consider only single-minded declarations from agents. Furthermore, from Lemma 5.6, all undominated strategies for agent i involve selecting a single set S_i and making a single-minded declaration for set S_i at its true value $v_i(S_i)$. An action in a given round therefore corresponds to a single set of items upon which to bid.

Of course, this still leaves an exponentially large (in m) space of actions, in general. One should therefore consider the format in which agent valuations are represented. For instance, a natural way to express an agent valuation is in XOR format [Sandholm 2002; Nisan 2000], which is a collection of single-minded valuations (i.e., set-value pairs) with the semantics that the value of a set S will be the maximum specified value for any set contained in S . The sets in the XOR representation are called the *desired sets* of agent i . This is an especially natural representation for many greedy algorithms, which typically iterate over the desired sets.

For a valuation represented in XOR format, it suffices to consider only bids for desired sets. This is because any other bid, say for a set S , would be dominated by a bid for the desired set that determines the value of S . Thus, if valuations are represented in XOR format, it suffices to consider a space of actions that is polynomial in the input size. One can therefore apply FPL in an efficient manner, leading to the following conclusion.

THEOREM 6.2. *Computing a regret-minimizing for agent i in mechanism $\mathcal{M}_2^*(\mathcal{A})$ can be done in time, per round, that is polynomial in n and the size of the XOR representation of agent i 's valuation.*

7. REPEATED COMBINATORIAL AUCTIONS AND BEST-RESPONSE AGENTS

In this section we consider the problem of designing mechanisms for agents that apply myopic best-response strategies asynchronously. Declaration sequence $D = (d^0, d^1, \dots)$ is an instance of *response dynamics* if for all prefixes, and for all $1 \leq t \leq T$, profiles d^{t-1} and d^t differ on the declaration of at most one player. Response dynamics D is an instance of *best-response dynamics* if, whenever d^{t-1} and d^t differ on the declaration of agent i , d_i^t maximizes agent i 's utility given the declarations of the other bidders. That is, $d_i^t \in \operatorname{argmax}_d \{u_i(d, d_{-i}^{t-1})\}$. In our model, agents are chosen for update uniformly at random to make a best-response, one agent per round. We will also assume that if a bidder is chosen for update but cannot improve his utility, he will choose to maintain his previous strategy.

We begin our analysis of myopic bidders by considering mechanism $\mathcal{M}_2^*(\mathcal{A})$ from Section 5 for a given monotone greedy algorithm \mathcal{A} . One might ask whether or not this mechanism converges to equilibrium under best-response dynamics. A simple example shows that this is not the case: there are circumstances in which mechanism $\mathcal{M}_2^*(\mathcal{A})$ has probability 0 of ever converging to a pure Nash equilibrium via best-response dynamics, despite the existence of a pure equilibrium.

Example 7.1. Consider a combinatorial auction with 6 agents and 4 objects, say $\{a, b, c, d\}$, under the feasibility constraint that each agent can receive at most 2 items. Let \mathcal{A} be the greedy allocation rule that allocates sets greedily by value. We consider an input instance given by the following set of true values (where the value for a set not listed is taken to be the maximum over its subsets).

player	set	value
1	$\{a, b\}$	4
1	$\{d\}$	6
2	$\{a\}$	2
2	$\{b, c\}$	5
3	$\{c\}$	4
4	$\{d\}$	5

Suppose the auction is resolved by mechanism $\mathcal{M}_2^*(\mathcal{A})$, and agents apply best-response dynamics. Agents 3 and 4 are single-minded and always maximize their utility by declaring truthfully. Agents 1 and 2 each have a strategic choice to make when bidding: which of their two desired sets should they bid upon? Note that once this decision is made, the way to bid is determined by Lemma 5.6 (i.e. bid truthfully for the desired set). It can be verified that from each of the resulting 4 possible declaration profiles, some player has incentive to change declaration. Thus best-response dynamics need not converge to an equilibrium.

The non-convergence of best-response dynamics is not a technical artifact of the model. Since the cyclic behavior described above employs undominated strategies by the players, we view it as a reasonable outcome to expect in such an auction (especially if players are not willing or able to randomize). The example above motivates a study of the total¹¹ (or equivalently the *average*) social welfare of $\mathcal{M}_2^*(\mathcal{A})$, over many rounds of best-response dynamics. We conjecture that, on average, the best-response dynamics on mechanism $\mathcal{M}_2^*(\mathcal{A})$ obtains an approximation to the optimal social welfare that is within a constant factor of the approximation ratio of the original algorithm \mathcal{A} .

CONJECTURE 7.1. *If \mathcal{A} is a c -approximate monotone strongly loser independent allocation rule, then $\mathcal{M}_2^*(\mathcal{A})$ has $O(c)$ price of (myopic) sinking.*

We leave the resolution of Conjecture 7.1 as an open problem. More generally, although we believe mechanism $\mathcal{M}_2^*(\mathcal{A})$ is an appropriate mechanism, the underlying goal is to have some black box transformation that converts a c -approximate monotone strongly loser independent allocation rule into a mechanism with $O(c)$ price of myopic sinking. As partial progress, we will focus on two specific combinatorial auction settings: the general combinatorial auction problem, and combinatorial auctions under a cardinality restriction. For these two settings, we will construct alternative mechanisms and analyze the welfare they generate under best-response dynamics.

7.1. The Approach

Our bound on the price of total anarchy of $\mathcal{M}_2^*(\mathcal{A})$ in Theorem 6.1 and our price of anarchy bound in Theorem 5.1 rely on a particular insight: if the social welfare of an auction outcome is low relative to the optimal welfare, then there must exist some agent i for whom the optimal assignment has a low critical price. We used this to argue that an outcome with low welfare cannot occur at equilibrium, since this agent i could improve his utility by pursuing his allocation in the optimal assignment.

The difficulty when extending this intuition to asynchronous best-response dynamics is that even if an agent can improve his utility by attempting to win some set for which the critical price is low, it may be that he has no chance to do so because he is not chosen to update his bid. Since each agent can expect to update his bid only once in every n rounds, our concern is that an agent spends most rounds wishing to make a

¹¹While other measures are possible, such as the minimum welfare over the cycle, the total or average welfare seems to be more relevant to a mechanism and is consistent with the regret minimization measure in section 6.

Mechanism \mathcal{M}_{sCA} :**Input:** Declaration profile $\mathbf{d} = d_1, \dots, d_n$.

1. $d' \leftarrow \text{SIMPLIFY}(\mathbf{d})$, % say $d'_i = (S_i, v_i)$
2. $(T_1, \dots, T_n) \leftarrow \mathcal{A}_{sCA}(\mathbf{d}')$.
3. For each i such that $T_i \neq \emptyset$:
4. $I \leftarrow \{j : S_j \cap T_i \neq \emptyset\}$. % I is the set of bids that intersect S_i
5. $\tau_i \leftarrow 2 \sum_{j \in I} d_j(S_j)$.
6. If $d'_i(T_i) \leq \tau_i$, set $T_i \leftarrow \emptyset$, $\tau_i \leftarrow 0$.
7. Allocate T_1, \dots, T_n and charge τ_1, \dots, τ_n . % Critical prices.

Fig. 3. Mechanism \mathcal{M}_{sCA} , an implementation of greedy algorithm \mathcal{A}_{sCA} for the s -CA problem. This mechanism uses procedure SIMPLIFY from Figure 2 in Section 5.3.

utility-improving bid, but that this improvement happens to be unavailable whenever it is that agent's turn to update.

We address this difficulty in two steps. First, we modify our mechanism so that the social welfare is never much less than the sum of the bids of *all* players - even those that are not allocated their desired sets. We accomplish this by designing the mechanism so that a winning bid must be significantly larger than the sum of all conflicting bids. This implies that if agent i places a large bid on round t , then we can think of agent i as making a large contribution to the social welfare even if she does not win her bid. Second, we demonstrate that with high probability, in almost half of the rounds (or more), either agent i places a large bid or else the critical price for her optimal allocation is high. Thus, even though agent i can modify her declaration only very infrequently, her (possibly indirect) contribution to the social welfare will still be large for approximately half of the rounds.

7.2. A Mechanism for s -CAs

Consider the s -CA problem, which is a combinatorial auction in which the feasibility constraint requires that no agent can be allocated more than s objects. An algorithm that greedily assigns sets in descending order by value obtains an s -approximation. Call this algorithm \mathcal{A}_{sCA} . We will construct a mechanism \mathcal{M}_{sCA} based on \mathcal{A}_{sCA} ; it is described in Figure 3. This algorithm simplifies incoming bids (in the same way as $\mathcal{M}_2^*(\mathcal{A})$) and runs algorithm \mathcal{A}_{sCA} to find a potential allocation. However, an additional condition for inclusion in the solution is imposed: the value declared for a set must be at least twice the sum of all bids for intersecting sets. Potential allocations that satisfy this condition are allocated, and the mechanism charges critical prices (that is, the smallest value at which an agent would be allocated their set by \mathcal{M}_{sCA} , which is not necessarily the same as the critical price for \mathcal{A}_{sCA}).

We note that since our mechanism implements a monotone algorithm and charges critical prices, Lemma 5.6 implies that undominated strategies for agent i involve choosing a set S_i and making a truthful single-minded bid for S_i at value $v_i(S_i)$. We will therefore assume that agents bid in this manner. However, the agent still has a strategic decision regarding which set S_i to choose.

We begin our analysis with some notation. Suppose that \mathbf{d} is a declaration profile, where each d_i is single-minded for some set S_i . For any set T , define the set of bids *intersecting* agent i 's bid for T in \mathbf{d} to be $I_i(\mathbf{d}, T) = \{j : j \neq i, S_j \cap T \neq \emptyset\}$. We also define $L_i(\mathbf{d}, T) = \{j : j \in I_i(\mathbf{d}, T), d_j(S_j) < v_i(T)\}$ to be the set of *lower* intersecting bids. We recursively define the set of *ancestors*, $A_i(\mathbf{d}, T)$, for agent i with respect to T and \mathbf{d} ; it is the set of all lower intersecting bids, plus all ancestors of those lower intersecting

bids. That is,

$$A_i(\mathbf{d}, T) = L_i(\mathbf{d}, T) \cup \bigcup_{j \in L_i(\mathbf{d}, T)} A_j(\mathbf{d}, S_j).$$

We say that \mathbf{d} is *separated for agent i* if $d_i(S_i) \geq 2 \sum_{j \in L_i(\mathbf{d}, S_i)} d_j(S_j)$. Profile \mathbf{d} is *separated* if it is separated for every bidder. Since an agent gains positive utility in mechanism \mathcal{M}_{cSA} only if the declaration is separated for him, we draw the following conclusion.

LEMMA 7.2. *If a declaration profile \mathbf{d} is separated, then it remains separated after a step of the best-response dynamics for mechanism \mathcal{M}_{sCA} .*

PROOF. Suppose agent i is chosen to update his bid, say from d_i to \tilde{d}_i . Let $\tilde{\mathbf{d}} = (\tilde{d}_i, \mathbf{d}_{-i})$. If agent i cannot improve his utility then $\tilde{\mathbf{d}} = \mathbf{d}$, so $\tilde{\mathbf{d}}$ is separated as required. Suppose otherwise, so agent i changes the set upon which he bids from, say, S_i to \tilde{S}_i . Since $u_i(\tilde{d}_i, \mathbf{d}_{-i}) > 0$, it must be that $\tilde{d}_i(\tilde{S}_i) > 2 \sum_{j \in L_i(\tilde{\mathbf{d}}, \tilde{S}_i)} d_j(S_j)$. This implies that $\tilde{\mathbf{d}}$ is separated for agent i . It remains to verify that $\tilde{\mathbf{d}}$ is separated for each $j \neq i$. Since $\tilde{\mathbf{d}}$ is separated for agent i , it must be that $\tilde{d}_i(\tilde{S}_i) > d_j(S_j)$ for all j such that $\tilde{S}_i \cap S_j \neq \emptyset$. This then implies that $i \notin L_j(\tilde{\mathbf{d}}, S_j)$ for each $j \neq i$. Moreover, only the declaration of agent i changes between \mathbf{d} and $\tilde{\mathbf{d}}$, and hence $L_j(\tilde{\mathbf{d}}, S_j)$ and $L_j(\mathbf{d}, S_j)$ can differ only on whether they include i . We conclude that, for all $j \neq i$, either $L_j(\tilde{\mathbf{d}}, S_j) = L_j(\mathbf{d}, S_j)$ or $L_j(\tilde{\mathbf{d}}, S_j) = L_j(\mathbf{d}, S_j) \setminus \{i\}$. In either case, $L_j(\tilde{\mathbf{d}}, S_j) \subseteq L_j(\mathbf{d}, S_j)$, and hence $\sum_{k \in L_j(\tilde{\mathbf{d}}, S_j)} \tilde{d}_k(S_k) \leq \sum_{k \in L_j(\mathbf{d}, S_j)} d_k(S_k)$. Therefore, since \mathbf{d} is separated for all $j \neq i$, $\tilde{\mathbf{d}}$ must be separated for each $j \neq i$ as well. \square

Motivated by Lemma 7.2, we will focus on separated declaration profiles¹² for the remainder of this section. We will show that, for any separated declaration profile, \mathcal{M}_{sCA} extracts a constant fraction of the sum of all declared bids, as welfare.

LEMMA 7.3. *For all separated declarations \mathbf{d} , $SW_{\mathcal{M}_{sCA}}(\mathbf{d}) \geq \frac{1}{2} \sum_i d_i(S_i)$.*

PROOF. We first claim that, for each agent i ,

$$d_i(S_i) \geq \sum_{j \in A_i(\mathbf{d}, S_i)} d_j(S_j). \quad (11)$$

¹²More formally, we could assume an initial empty declaration (which is trivially separated) so that by induction all declarations will be separated. Alternatively, we can modify mechanism \mathcal{M}_{sCA} so that, with vanishingly small probability, an alternative allocation rule is used. This alternative rule chooses an agent at random, and assigns him all objects at no cost *as long as the input declaration is separated for that agent*. Thus, any separated declaration by agent i results in positive expected utility. Since any non-separated declaration by an agent results in a utility of 0 for that agent, it must be that the utility-maximizing declaration by any agent must be separated.

This follows by structural induction on the recursive definition of $A_i(\mathbf{d}, T)$, since for each i we have

$$\begin{aligned} d_i(S_i) &\geq 2 \sum_{j \in L_i(\mathbf{d}, S_i)} d_j(S_j) \\ &\geq \sum_{j \in L_i(\mathbf{d}, S_i)} d_j(S_j) + \sum_{j \in L_i(\mathbf{d}, S_i)} \sum_{k \in A_j(\mathbf{d}, S_j)} d_k(S_k) \\ &\geq \sum_{j \in A_i(\mathbf{d}, S_i)} d_j(S_j) \end{aligned}$$

where the first inequality follows by separatedness.

Let $N \subseteq [n]$ be the set of agents that receive non-empty sets in $\mathcal{M}_{sCA}(\mathbf{d})$. We next claim that for all $j \notin N$, either $S_j = \emptyset$ or there exists some $i \in N$ such that $j \in A_i(\mathbf{d}, S_i)$. To see this, let i be such that $d_i(S_i)$ is maximal, subject to $j \in A_i(\mathbf{d}, S_i)$. Note that such an i must exist whenever $S_j \neq \emptyset$, though it could be that $i = j$. By maximality, the bid of i cannot intersect with any larger bid in \mathbf{d} , so i is allocated a nonempty set by $\mathcal{A}_{sCA}(\mathbf{d})$. Moreover, by separatedness, the allocation to i is not set to \emptyset on line 6 of $\mathcal{M}_{sCA}(\mathbf{d})$. Thus $i \in N$, as claimed.

Now, taking a sum over all $j \notin N$ and applying (11), we can conclude

$$\sum_{j \notin N} d_j(S_j) \leq \sum_{i \in N} \sum_{j \in A_i(\mathbf{d}, S_i)} d_j(S_j) \leq \sum_{i \in N} d_i(S_i). \quad (12)$$

We then have

$$\begin{aligned} \sum_i d_i(S_i) &= \sum_{i \in N} d_i(S_i) + \sum_{j \notin N} d_j(S_j) \\ &\leq \sum_{i \in N} d_i(S_i) + \sum_{i \in N} d_i(S_i) \\ &= 2SW_{\mathcal{M}_{sCA}}(\mathbf{d}) \end{aligned}$$

as required. \square

We are now ready to bound the price of sinking for \mathcal{M}_{sCA} . We will achieve the following bound.

THEOREM 7.4. *Choose $\epsilon > 0$ and suppose $D = d^1, \dots, d^T$ is an instance of best-response dynamics with random player order, where agents play undominated strategies, and $T > \epsilon^{-1}n$. Then*

$$SW_{\mathcal{M}_{sCA}}(D) \geq \left(\frac{1 - 2\epsilon}{16s + 8} \right) SW_{OPT}(\mathbf{v})$$

with probability at least $1 - ne^{-T\epsilon^2/32n}$.

Before proving Theorem 7.4, let us make some remarks. If we take ϵ to be a small constant and assume $T = \Omega(n^{1+\delta})$ for some $\delta > 0$, we conclude that $SW_{\mathcal{M}_{sCA}}(D) > \frac{1}{O(s)} SW_{opt}(\mathbf{v})$ with high probability. Thus \mathcal{M}_{sCA} implements an $O(s)$ approximation to the s -CA problem for best-response bidders over sufficiently many rounds. In other words, \mathcal{M}_{sCA} has $O(s)$ price of (myopic) sinking.

We now begin with the proof of the theorem. Let \mathbf{y} be an optimal allocation with respect to the agents' true types \mathbf{v} . For a given time step $t \leq T$, we will define a notion of "good" agents on step t . Let G_1^t denote the set of agents for which $d_i^t(S_i) \geq \frac{1}{2}v_i(y_i)$, and let G_2 denote the set of agents for which $\sum_{j \in I_i(\mathbf{d}^t, y_i)} d_j^t(S_j) > \frac{1}{4}v_i(y_i)$. That is, G_1^t

is the set of agents making relatively large bids at time t , and G_2^t is the set of agents i for whom a relatively large bid (that is, up to $\frac{1}{4}v_i(y_i)$) on set y_i would not be a winning bid. Write $G^t = G_1^t \cup G_2^t$; we will refer to G^t as the set of good agents at time t .

We will first argue that, at any given time t , the welfare achieved by our mechanism achieves a good approximation to the welfare of the optimal outcome, restricted to agents in G^t .

LEMMA 7.5. *For all t , $SW_{\mathcal{M}_{sCA}}(\mathbf{d}^t) \geq \frac{1}{8s+4} \sum_{i \in G^t} v_i(y_i)$.*

PROOF. By Lemma 7.3 and the definition of G_1^t , we have

$$SW_{\mathcal{M}_{sCA}}(\mathbf{d}^t) \geq \frac{1}{2} \sum_i d_i^t(S_i) \geq \frac{1}{2} \sum_{i \in G_1^t} d_i^t(S_i) \geq \frac{1}{4} \sum_{i \in G_1^t} v_i(y_i). \quad (13)$$

Next, by Lemma 7.3 and the definition of G_2^t , we have

$$\begin{aligned} SW_{\mathcal{M}_{sCA}}(\mathbf{d}^t) &\geq \frac{1}{2} \sum_i d_i^t(S_i) \\ &\geq \frac{1}{2} \cdot \frac{1}{s} \left(\sum_i \sum_{j \in I_i(\mathbf{d}^t, y_i)} d_j^t(S_j) \right) \\ &\geq \frac{1}{2s} \sum_{i \in G_2^t} \sum_{j \in I_i(\mathbf{d}^t, y_i)} d_j^t(S_j) \\ &\geq \frac{1}{8s} \sum_{i \in G_2^t} v_i(y_i) \end{aligned} \quad (14)$$

where the second inequality follows because each set S_j can intersect at most $|S_j| \leq s$ sets y_i , and the final inequality is by the definition of G_2^t . Combining (13) and (14) yields the desired result. \square

In light of Lemma 7.5, our goal is to show that each agent is good sufficiently often over the course of an instance D of best-response dynamics. To this end, we will establish bounds on the probability that an agent lies in G_1 or G_2 .

LEMMA 7.6. *For any i and t , $\Pr[i \in G_1^{t+1} | i \in G_1^t] \geq 1 - \frac{1}{n}$.*

PROOF. If $i \in G_1^t$, and i is not selected by the best-response dynamics following round t , then it must be that $i \in G_1^{t+1}$. Since i is selected with probability $1/n$, the result follows. \square

LEMMA 7.7. *For any i and t , $\Pr[i \in G_1^{t+1} | i \notin G^t] \geq \frac{1}{n}$.*

PROOF. Suppose $i \notin G^t$, and that i is selected by the best-response dynamics process following round t . Since $i \notin G_2^t$, it must be that that $\theta_i(y_i, \mathbf{d}_{-i}) = 2 \sum_{j \in I_i(\mathbf{d}, y_i)} d_j(S_j)$ in mechanism \mathcal{M}_{sCA} , so agent i would obtain utility at least $\frac{1}{2}v_i(y_i)$ by making a single-minded declaration for set y_i at value $v_i(y_i)$. His utility-maximizing declaration must therefore make at least this much utility, and therefore is a bid for some set S_i with $v_i(S_i) \geq \frac{1}{2}v_i(y_i)$. Thus, with probability $\frac{1}{n}$, agent i is selected and necessarily chooses d_i^{t+1} such that $i \in G_1^{t+1}$. \square

We now show that the above observations imply that any given agent will be in G^t reasonably often, with high probability. This will follow immediately from the following technical lemma, whose proof we defer:

LEMMA 7.8. *Suppose that $\{A^t\}_{t \leq T}$ and $\{B^t\}_{t \leq T}$ are sequences of binary random variables satisfying the following properties:*

- (1) $A^t + B^t \leq 1$ for all t ,
- (2) $\Pr[A^{t+1} = 1 | B^t = 1] \geq \frac{1}{n}$ for all t , and
- (3) $\Pr[A^{t+1} = 1 | A^t = 1] \geq 1 - \frac{1}{n}$ for all t .

Then $\Pr[\sum_t B^t \geq (\frac{1}{2} + \epsilon)T] \leq e^{-T\epsilon^2/32n}$.

COROLLARY 7.9. *For any agent i , with probability at least $1 - e^{-T\epsilon^2/32n}$, agent i will be in G^t for at least $(\frac{1}{2} - \epsilon)T$ values of t .*

PROOF. Let A^t be the indicator for event $i \in G_1^t$, and let B^t be the indicator for the event $i \notin G^t$. By Lemma 7.6 and Lemma 7.7, we can apply Lemma 7.8 to conclude that $i \notin G^t$ for at most $(\frac{1}{2} + \epsilon)T$ values of t , with the required probability. \square

We are now ready to prove the main result of this section.

Proof of Theorem 7.4 : Corollary 7.9 implies that each agent i will be in G^t for at least $(\frac{1}{2} - \epsilon)T$ values of t , with probability at least $1 - e^{-T\epsilon^2/32n}$. The union bound then implies that this occurs for every agent with probability at least $1 - ne^{-T\epsilon^2/32n}$. Conditioning on the occurrence of this event, Lemma 7.5 implies

$$\begin{aligned} SW_{\mathcal{M}_{sCA}}(D) &= \frac{1}{T} \sum_t SW_{\mathcal{M}_{sCA}}(\mathbf{d}^t) \\ &\geq \frac{1}{(8s+4)T} \sum_t \sum_{i \in G^t} v_i(y_i) \\ &\geq \frac{1}{(8s+4)T} \sum_i T \left(\frac{1}{2} - \epsilon\right) v_i(y_i) \\ &\geq \left(\frac{1-2\epsilon}{16s+8}\right) SW_{OPT}(\mathbf{v}) \end{aligned}$$

which implies the required bound. \square

7.3. Proof of Lemma 7.8

Our proof will make use of the *method of average bounded differences*. We will begin by giving a brief statement of this technique; see, for example, [Dubhashi and Panconesi 2009] for a more thorough treatment. Suppose that z_1, \dots, z_n are (not necessarily independent) random variables, and let f be real-valued function of z_1, \dots, z_n satisfying the property that, for each $i \in [n]$ and any two values a, a' that z_i can assume, there is a non-negative value c_i such that

$$|E[f|z_1, \dots, z_{i-1}; z_i = a] - E[f|z_1, \dots, z_{i-1}; z_i = a']| \leq c_i$$

where the expectations are with respect to the values of z_{i+1}, \dots, z_n . Then the method of average bounded differences states that $\Pr[f > \mathbf{E}[f] + \ell] \leq e^{-\ell^2/2c}$ for all $\ell > 0$, where $c = \sum_{i \in [n]} c_i^2$.

We now proceed with the proof of Lemma 7.8. First recall the statement of the lemma. Suppose that $\{A^t\}_{t \leq T}$ and $\{B^t\}_{t \leq T}$ are sequences of binary random variables satisfying the following properties:

- (1) $A^t + B^t \leq 1$ for all t ,
- (2) $\Pr[A^{t+1} = 1 | B^t = 1] \geq \frac{1}{n}$ for all t , and

(3) $\Pr[A^{t+1} = 1 | A^t = 1] \geq 1 - \frac{1}{n}$ for all t .

We wish to prove that $\Pr[\sum_t B^t \geq (\frac{1}{2} + \epsilon)T] \leq e^{-T\epsilon^2/32n}$.

Consider the steps in which A^t and B^t are not both 0; let $R = \{t: A^t + B^t \geq 1\}$ be this set of steps. For all $r \leq |R|$, let $t(r)$ denote the r th largest element of R . That is, $t(r)$ is the step at which either A^t or B^t is 1 for the r th time. Let C^r denote the (indicator variable for the) event that on the r th step in which either A^t or B^t is 1, it is B^t that is 1. Then $\sum_{r \leq T} C_r = \sum_{t \leq T} B^t$, so it is enough to show that $\Pr[\sum_{r \leq T} C_r > (\frac{1}{2} + \epsilon)T] < e^{-T\epsilon^2/32n}$.

Since $\Pr[A^{t+1} = 1 | B^t = 1] \geq 1/n$, we have that $\Pr[C_{r+1} | C_r] \leq (1 - 1/n)$. Moreover, if C_r does not occur, then $A^{t(r)}$ occurs, and hence (since $\Pr[A^{t+1} = 1 | A^t = 1] \geq 1 - \frac{1}{n}$), C_{r+1} occurs with probability at most $1/n$. That is, $\Pr[C_{r+1} | \neg C_r] \leq 1/n$.

Let D_1, D_2, \dots, D_T be a random walk on $\{0, 1\}$ defined by $\Pr[D_r | D_{r-1}] = (1 - 1/n)$, $\Pr[D_r | \neg D_{r-1}] = 1/n$, and initial condition D_0 . Then our bounds above imply that $\sum_r C_r$ is stochastically dominated by $\sum_r D_r$, and hence $\Pr[\sum_r C_r > (\frac{1}{2} + \epsilon)T] \leq \Pr[\sum_r D_r > (\frac{1}{2} + \epsilon)T]$. It will therefore suffice to show that $\Pr[\sum_{r \leq T} D_r > (\frac{1}{2} + \epsilon)T] < e^{-T\epsilon^2/32n}$.

The definition of D_r yields

$$\begin{aligned} \Pr[D_r] &= \frac{1}{n}(1 - \Pr[D_{r-1}]) + \left(1 - \frac{1}{n}\right) \Pr[D_{r-1}] \\ &= \frac{1}{n} + \left(1 - \frac{2}{n}\right) \Pr[D_r]. \end{aligned}$$

Solving this linear recurrence (with initial condition $D_0 \in \{0, 1\}$) yields

$$\begin{aligned} \Pr[D_r] &= \frac{1}{2}(1 - (1 - 2/n)^r) + D_0(1 - 2/n)^r \\ &= \frac{1}{2} + (D_0 - \frac{1}{2})(1 - 2/n)^r. \end{aligned}$$

Linearity of expectation then implies

$$E\left[\sum_r D_r\right] = \frac{1}{2}T + \left(D_0 - \frac{1}{2}\right) \frac{n}{2}(1 - (1 - 2/n)^{T+1}).$$

From this we conclude that

$$E\left[\sum_r D_r\right] < \frac{1}{2}T + \frac{n}{4} \tag{15}$$

and moreover

$$\left|E\left[\sum_r D_r \mid D_0 = 1\right] - E\left[\sum_r D_r \mid D_0 = 0\right]\right| < \frac{n}{2}. \tag{16}$$

Let $k = T/n$ and define random variables F_1, \dots, F_k by $F_i = \sum_{r \in [in, (i+1)n-1]} D_r$. Note that $F_i \in [0, n]$ for all i . Note also that

$$\sum_r D_r = \sum_i F_i. \tag{17}$$

We would now like to apply the method of average bounded differences to random variables F_1, \dots, F_k and function $f = \sum_i F_i$. To do so, we must consider the expectation

$\mathbf{E}[\sum_i F_i | F_1, \dots, F_{i-1}, F_i = \alpha]$ for $\alpha \in [0, n]$. But note that the influence of F_ℓ on the values of $F_{\ell+1}, \dots, F_k$ is captured entirely by the value of $D_{(\ell+1)n-1}$, and from (16) the influence of $D_{(\ell+1)n-1}$ on $\sum_{r=(\ell+1)n}^T D_r = \sum_{r=\ell+1}^k F_r$ is bounded by $\frac{n}{2}$. Since the value of F_ℓ also influences the sum $\sum_i F_i$ directly by at most n (due to its being included in the summation), we conclude that for all $\alpha, \alpha' \in [0, n]$,

$$\left| \mathbf{E} \left[\sum_j F_j \mid F_1, \dots, F_{i-1}, F_i = \alpha \right] - \mathbf{E} \left[\sum_j F_j \mid F_1, \dots, F_{i-1}, F_i = \alpha' \right] \right| \leq 3n/2.$$

Thus, by the method of average bounded differences (and recalling that $k = T/n$), we conclude that

$$\Pr \left[\sum_j F_j > \mathbf{E} \left[\sum_j F_j \right] + (\epsilon/2)T \right] \leq e^{-(T\epsilon/2)^2/2(3n/2)^2k} < e^{-T\epsilon^2/18n}. \quad (18)$$

Our final step is to bound $\mathbf{E} \left[\sum_j F_j \right] + (\epsilon/2)T$ from the left-hand side of (18). Since $T > \epsilon^{-1}n$, we have from (17) and (15) that

$$\mathbf{E} \left[\sum_j F_j \right] + (\epsilon/2)T = \mathbf{E} \left[\sum_j D_j \right] + (\epsilon/2)T \leq \frac{1}{2}T + \frac{n}{4} + (\epsilon/2)T < \left(\frac{1}{2} + \epsilon \right) T.$$

Thus (18) implies

$$\Pr \left[\sum_j F_j > \left(\frac{1}{2} + \epsilon \right) T \right] < e^{-T\epsilon^2/32n}$$

and the result follows.

7.4. A Mechanism for General CAs

Consider the following algorithm for the general CA problem: first try greedily (by value) assigning sets of size at most \sqrt{m} , then try allocating all items to the single agent with the highest declared value; return whichever of those two solutions generates more welfare. This algorithm is known to be an $O(\sqrt{m})$ approximation for the general CA problem [Halldórsson 2000]. We will construct a mechanism \mathcal{M}_{rCA} based on this algorithm; it is described in Figure 5. The mechanism \mathcal{M}_{CA} essentially implements two copies of \mathcal{M}_{sCA} (as described in the previous section): one for sets of size at most \sqrt{m} (which we will call $\mathcal{M}_{\sqrt{m}CA}$), and one for allocating all objects to a single bidder. We call the latter the “grand bundle mechanism” \mathcal{M}_{GB} , given in Figure 4. It is tempting to simply deterministically return whichever of the two solutions yields the largest social welfare. However, we will see in Appendix B that this deterministic algorithm will not always have a good POA in the one shot game. Instead we will consider a randomized mechanism in each round, one that randomly chooses between $\mathcal{M}_{\sqrt{m}CA}$ and \mathcal{M}_{GB} .

The analysis of the average social welfare obtained by \mathcal{M}_{CA} closely follows the analysis for \mathcal{M}_{sCA} . Our high-level approach is to apply this analysis twice: once for allocations of sets of size at most \sqrt{m} , and once for allocations of all objects to a single bidder. The final result is the following.

THEOREM 7.10. *Choose $\epsilon > 0$ and suppose $D = d^1, \dots, d^T$ is an instance of best-response dynamics with random player order, where agents play undominated strate-*

Mechanism \mathcal{M}_{GB} :**Input:** Declaration profile $\mathbf{d} = d_1, \dots, d_n$.% A Vickrey auction, thinking of M as a single item.

1. Let $(T_1, \dots, T_n) \leftarrow (\emptyset, \dots, \emptyset)$.
2. Let $j \leftarrow \operatorname{argmax}_j \{d_j(M)\}$.
3. Set $T_j \leftarrow M$.
4. Return (T_1, \dots, T_n) and charge critical prices.

Fig. 4. Mechanism \mathcal{M}_{GB} , an implementation of the grand bundle algorithm.**Mechanism \mathcal{M}_{rCA} :****Input:** Declaration profile $\mathbf{d} = d_1, \dots, d_n$.0 For each i , define d_i' to be $d_i'(T) = \max_{S \subseteq T, |S| \leq \sqrt{m}} \{d_i(S)\}$.

1. With probability $1/2$, return $\mathcal{M}_{sCA}(\mathbf{d}')$ for $s = \sqrt{m}$.
2. Else return $\mathcal{M}_{GB}(\mathbf{d})$.

Fig. 5. Mechanism \mathcal{M}_{rCA} , randomized mechanism for the CA problem.

gies, and $T > \epsilon^{-1}n$. Then

$$SW_{\mathcal{M}_{rCA}}(D) \geq \left(\frac{1 - 2\epsilon}{O(\sqrt{m})} \right) SW_{opt}(\mathbf{v})$$

with probability at least $1 - 2ne^{-T\epsilon^2/32n}$.

PROOF. The initializations in mechanisms \mathcal{M}_{GB} , \mathcal{M}_{sCA} , and \mathcal{M}_{rCA} ensure that declared valuations are used in the following way: with probability $1/2$, only the bids on M are used; otherwise, the mechanism considers only the single highest bid on any set of size at most \sqrt{m} . Given this, it is without loss of generality to assume that each declaration d_i is the XOR of two single-minded bids: one for the set M of all items, and one for a set S_i of size at most \sqrt{m} .

Since \mathcal{M}_{GB} is a Vickrey auction on the grand bundle, undominated strategies for \mathcal{M}_{GB} involve bidding truthfully on M . Also, our prior analysis of \mathcal{M}_{sCA} implies that undominated strategies for \mathcal{M}_{sCA} involve bidding truthfully on the chosen set S_i . Thus, for mechanism \mathcal{M}_{rCA} the undominated strategies are precisely those that bid truthfully both on M and on S_i .

Let \mathbf{y} be the optimal allocation in which each non-empty set has size greater than \sqrt{m} . Then $SW(\mathbf{y}, \mathbf{v}) \leq \sqrt{m} \max_i v_i(M)$, since there can be at most \sqrt{m} non-empty sets in \mathbf{y} . Since we can assume the agent will be bidding truthfully for set M in \mathcal{M}_{GB} (and has no strategic decision as to which set to bid for as in \mathcal{M}_{sCA}), we can immediately apply the above \sqrt{m} approximation analysis. Recalling that we are calling mechanism \mathcal{M}_{GB} with probability $\frac{1}{2}$ in \mathcal{M}_{rCA} , we conclude that

$$SW_{\mathcal{M}_{rCA}}(D) \geq \frac{1}{2} \cdot \frac{1}{\sqrt{m}} SW(\mathbf{y}, \mathbf{v}).$$

Now let \mathbf{z} be the optimal allocation of sets of size at most \sqrt{m} . From Theorem 7.4 and again recalling that we are calling $\mathcal{M}_{\sqrt{m}CA}$ with probability $\frac{1}{2}$, we conclude that

with probability at least $1 - ne^{-T\epsilon^2/32n}$,

$$SW_{\mathcal{M}_{rCA}}(D) \geq \frac{1}{2} \cdot \frac{1 - 2\epsilon}{8\sqrt{m}} SW(\mathbf{z}, \mathbf{v}).$$

Taking the union bound over the events described above, and noting that $SW_{OPT}(\mathbf{v}) \leq SW(\mathbf{z}, \mathbf{v}) + SW(\mathbf{y}, \mathbf{v})$, we conclude that, with probability at least $1 - ne^{-T\epsilon^2/32n}$,

$$SW_{\mathcal{M}_{rCA}}(D) \geq \frac{1 - 2\epsilon}{18\sqrt{m}} SW_{OPT}(\mathbf{v})$$

as required.

□

We conclude that mechanism \mathcal{M}_{rCA} implements an $O(\sqrt{m})$ approximation to the combinatorial auction problem for best-response bidders, with high probability, whenever $T = \Omega(n^{1+\delta})$ for $\delta > 0$.

8. CONCLUSION AND OPEN PROBLEMS

A central theme in algorithmic mechanism design concerns the transformation of algorithms into mechanisms that satisfy some game-theoretic solution concept (e.g incentive compatibility, approximations at equilibrium). In contrast to incentive compatibility (where generally we do not expect to be able to preserve approximation bounds), we show that for a wide class of greedy algorithms, approximation bounds for combinatorial allocation algorithms can be transformed into mechanisms that enjoy closely matching price of anarchy bounds. Notably, these results apply to Bayesian equilibria and some forms of repeated auctions.

We leave open a number of interesting challenges. Our results are motivated by, and pertain to, monotone greedy algorithms as formally defined in Section 2.4. In fact, the key property of such algorithms are that they are monotone strongly loser-independent as defined in Section 3 and, with the exception of the results for correlated equilibria and best response dynamics, our results hold for arbitrary monotone strongly loser-independent algorithms. In particular, our result for correlated equilibria of the first-price mechanism requires that the allocation algorithm \mathcal{A} is a fixed order greedy algorithm and achieves a price of anarchy bound of $4c$, in contrast to our $c + o(1)$ result for independent agent distributions. Can the price of anarchy bound for correlated equilibria be improved? Can it be extended to adaptive greedy or more generally strongly loser independent algorithms? Our results for best response dynamics are restricted to particular greedy algorithms for the combinatorial auction problem and we lack a general approach that will work for all greedy allocation algorithms.

Greedy algorithms for allocation problems often provide the best known approximations for combinatorial auction problems, but are nevertheless a restricted class of algorithms. The basic open question in this regard is: for what class of allocation algorithms can a given approximation algorithm \mathcal{A} be transformed into a deterministic or randomized mechanism $\mathcal{M}(\mathcal{A})$ that provides a POA bound (closely) matching \mathcal{A} 's approximation ratio? We also note that our framework does not capture *all* algorithms that are typically thought of as greedy, since our definition assumes that it is the player-allocation pairs that are considered greedily. This excludes, for example, the greedy algorithm for combinatorial auctions where the valuation function of every agent is a monotone submodular function. That algorithm considers each item (in any arbitrary order) and awards it to the agent having the maximum marginal gain for that item. This suggests the question as to whether or not price of anarchy re-

sults could be extended to more general forms of greedy allocation rules. Similarly, the recent Buchbinder et al [Buchbinder et al. 2012] randomized online algorithm for unconstrained non-monotone submodular maximization also considers items (rather than bids) in a greedy algorithm. Can our methodology be extended to include non monotone combinatorial auctions (i.e. no free disposal)? It is also interesting to consider more general settings of incomplete information, such as interdependent valuations — See Roughgarden and Talgam-Cohen [Roughgarden and Talgam-Cohen 2013].

REFERENCES

- A. Archer and E. Tardos. 2001. Truthful mechanisms for one-parameter agents. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*.
- M. Babaioff and L. Blumrosen. 2004. Computationally-Feasible Truthful Auctions for Convex Bundles. In *Proc. 7th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*.
- P. Baptiste. 1999. Polynomial time algorithms for minimizing the weighted number of late jobs on a single machine with equal processing times. *Journal of Scheduling* 2 (1999), 245–252. Issue 6.
- Y. Bartal, R. Gonen, and N. Nisan. 2003. Incentive Compatible Multi Unit Combinatorial Auctions. In *Proc. 9th Conf. on Theoretical Aspects of Rationality and Knowledge*.
- K. Bhawalkar and T. Roughgarden. 2011. Welfare Guarantees for Combinatorial Auctions with Item Bidding. In *Proc. 22nd ACM Symp. on Discrete Algorithms*.
- A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. 2008. Regret Minimization and the Price of Total Anarchy. In *Proc. 39th ACM Symp. on Theory of Computing*.
- Avrim Blum and Yishay Mansour. 2007. From External to Internal Regret. *Journal of Machine Learning Research* 8 (2007), 1307–1324.
- A. Borodin and B. Lucier. 2010. Greedy Mechanism Design for Truthful Combinatorial Auctions. In *Proc. 37th Intl. Colloq. on Automata, Languages and Programming*.
- A. Borodin, M. N. Nielsen, and C. Rackoff. 2002. (Incremental) Priority Algorithms. In *Proc. 13th ACM Symp. on Discrete Algorithms*.
- P. Briest, P. Krysta, and B. Vöcking. 2005. Approximation Techniques for Utilitarian Mechanism Design. In *Proc. 36th ACM Symp. on Theory of Computing*.
- Niv Buchbinder, Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. 2012. A Tight Linear Time (1/2)-Approximation for Unconstrained Submodular Maximization. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS '12)*. IEEE Computer Society, Washington, DC, USA, 649–658. DOI: <http://dx.doi.org/10.1109/FOCS.2012.73>
- Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos. 2012. On the efficiency of equilibria in generalized second price auctions. *CoRR* abs/1201.6429 (2012).
- C. Chekuri and I. Gamzu. 2009. Truthful Mechanisms via Greedy Iterative Packing. In *Proc. 12th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*.
- G. Christodoulou, A. Kovács, and Michael Schapira. 2008. Bayesian Combinatorial Auctions. In *Proc. 35th Intl. Colloq. on Automata, Languages and Programming*. 820–832.
- P. Cramton, Y. Shoham, and R. Steinberg (eds.). 2005. *Combinatorial Auctions*. MIT Press, Cambridge, MA.
- Bart de Keijzer, Evangelos Markakis, Guido Schfer, and Orestis Telelis. 2013. Inefficiency of Standard Multi-unit Auctions. In *Algorithms ESA 2013*, HansL. Bodlaender and GiuseppeF. Italiano (Eds.). Lecture Notes in Computer Science, Vol. 8125. Springer Berlin Heidelberg, 385–396. DOI: http://dx.doi.org/10.1007/978-3-642-40450-4_33
- S. Dobzinski. 2011. An Impossibility Result for Truthful Combinatorial Auctions with Submodular Valuations. In *Proc. 42nd ACM Symp. on Theory of Computing*.
- Shahar Dobzinski and Jan Vondrak. 2012. The Computational Complexity of Truthfulness in Combinatorial Auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12)*. ACM, New York, NY, USA, 405–422. DOI: <http://dx.doi.org/10.1145/2229012.2229044>
- D. Dubhashi and A. Panconesi. 2009. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press.
- Shaddin Dughmi and Jan Vondrák. 2011. Limitations of Randomized Mechanisms for Combinatorial Auctions. In *FOCS*. 502–511.

- B. Edelman, M. Ostrovsky, and M. Schwarz. 2005. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. In *Stanford Graduate School of Business Research Paper No. 1917*.
- Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. 2013. Simultaneous Auctions Are (Almost) Efficient. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing (STOC '13)*. New York, NY, USA, 201–210.
- M. Goemans, V. Mirrokni, and A. Vetta. 2005. Sink Equilibria and Convergence. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*.
- Magnús M. Halldórsson. 2000. Approximations of Weighted Independent Set and Hereditary Subset Problems. *J. Graph Algorithms Appl.* 4, 1 (2000).
- J. Hannan. 1957. Approximation to Bayes risk in repeated plays. In *Contributions to the Theory of Games*, M. Dresher, A. Tucker, and P. Wolfe (Eds.), Vol. 4. Princeton University Press.
- Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. 2011. Non-price equilibria in markets of discrete goods. In *EC*. 295–296.
- A. Kalai and S. Vempala. 2005. Efficient algorithms for online decision problems. *J. Comput. System Sci.* 71, 3 (2005), 291 – 307. DOI : <http://dx.doi.org/DOI:10.1016/j.jcss.2004.10.016> Learning Theory 2003.
- D. Lehmann, L. I. O’Callaghan, and Y. Shoham. 1999. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. 1st ACM Conf. on Electronic Commerce*. ACM Press, 96–102.
- B. Lucier. 2010. Beyond Equilibria: Mechanisms for Repeated Combinatorial Auctions. In *Proc. 1st Symp. on Innovations in Computer Science*.
- B. Lucier. 2011. The Power of Uncertainty: Algorithmic Mechanism Design in Settings of Incomplete Information. (2011). University of Toronto, PHD thesis.
- B. Lucier and A. Borodin. 2010. Price of Anarchy for Greedy Auctions. In *Proc. 21st ACM Symp. on Discrete Algorithms*.
- Brendan Lucier and Renato Paes Leme. 2011. GSP auctions with correlated types. In *EC*. 71–80.
- Evangelos Markakis and Orestis Telelis. 2012. Uniform Price Auctions: Equilibria and Efficiency. In *Algorithmic Game Theory*, Maria Serna (Ed.). Springer Berlin Heidelberg, 227–238. DOI : http://dx.doi.org/10.1007/978-3-642-33996-7_20
- A. Mu’alem and N. Nisan. 2008. Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. *Games and Economic Behavior* 64 (2008), 612–631. Issue 2.
- N. Nisan. 2000. Bidding and allocation in combinatorial auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce*.
- N. Nisan and A. Ronen. 1999. Algorithmic Mechanism Design. In *Proc. 31st ACM Symp. on Theory of Computing*. ACM Press, 129–140.
- C. Papadimitriou. 2001. Algorithms, Games and the Internet. In *Proc. 33rd ACM Symp. on Theory of Computing*. ACM Press, 749–752.
- C. Papadimitriou, M. Schapira, and Y. Singer. 2008. On the Hardness of Being Truthful. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*.
- T. Roughgarden. 2009. Intrinsic Robustness of the Price of Anarchy. In *Proc. 40th ACM Symp. on Theory of Computing*.
- Tim Roughgarden. 2012. Intrinsic robustness of the price of anarchy. *Commun. ACM* 55, 7 (2012), 116–123.
- Tim Roughgarden and Inbal Talgam-Cohen. 2013. Optimal and near-optimal mechanism design with interdependent values. In *EC*. 767–784.
- T. Roughgarden and Eva Tardos. 2000. How Bad is Selfish Routing?. In *41st Annual Symposium on Foundations of Computer Science*. 93–102.
- T. Sandholm. 2002. Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence* 135 (2002), 1–54.
- Vasilis Syrgkanis and Éva Tardos. 2013. Composable and efficient mechanisms. In *STOC*. 211–220.

A. EXISTENCE OF PURE NASH EQUILIBRIA

As stated in section 4, the power of our pure price of anarchy bounds such as Theorem 4.3, 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 for $\mathcal{M}_1(\mathcal{A})$ is given below.

Example A.1. Consider an instance of the combinatorial auction problem with two objects, $M = \{a, b\}$, and three agents. Our feasibility constraint is that 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 greedy algorithm that ranks bids by value. Suppose the true types of the agents are as follows: $v_1(a) = 4$, $v_1(b) = 2$, $v_2(a) = 3$, $v_2(b) = 0$, $v_3(a) = 0$, and $v_3(b) = 3$.

We now prove that no pure Nash equilibrium exists for this example, even if we assume that agents declare multiples¹³ of some $\epsilon > 0$. Assume for contradiction that there is a Nash equilibrium \mathbf{d} for type profile \mathbf{v} 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 (so he would certainly 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.

B. 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 [Bartal et al. 2003; Briest et al. 2005; Halldórsson 2000; Mu’alem and Nisan 2008]. When bidders are single-minded, such a combination rule will be incentive-compatible [Mu’alem and Nisan 2008]. 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 B.1. 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 [Mu’alem and Nisan 2008].

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 , v_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{v} 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

¹³That is, our lack of pure equilibrium is not due to the possibility of infinitesimal improvements. One can also interpret our example as demonstrating that there is no $(1 + \epsilon)$ -approximate pure Nash equilibrium for small $\epsilon > 0$.

agent, for a total welfare of $1 + \epsilon$. Also, \mathbf{d} is a pure Nash equilibrium for $\mathcal{M}_1(\mathcal{A}_{max})$ and $\mathcal{M}_{crit}(\mathcal{A}_{max})$: 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 example above, one must consider different ways to combine two allocation rules. For instance, one could implement each rule as a separate mechanism, then randomly choose between the two with equal probability. This is the approach taken in Section 7.4. Such an approach can work well when the two allocation rules work with disjoint parts of the declaration space, so that agents can optimize their bids separately for each mechanism.