

Beyond Equilibria: Mechanisms for Repeated Combinatorial Auctions

Brendan Lucier¹

¹Dept of Computer Science, University of Toronto
blucier@cs.toronto.edu

Abstract: We study the design of mechanisms in combinatorial auction domains. We focus on settings where the auction is repeated, motivated by auctions for licenses or advertising space. We consider models of agent behaviour in which they either apply common learning techniques to minimize the regret of their bidding strategies, or apply short-sighted best-response strategies. We ask: when can a black-box approximation algorithm for the base auction problem be converted into a mechanism that approximately preserves the original algorithm’s approximation factor on average over many iterations? We present a general reduction for a broad class of algorithms when agents minimize external regret. We also present a mechanism for the combinatorial auction problem that attains an $O(\sqrt{m})$ approximation on average when agents apply best-response dynamics.

Keywords: Combinatorial auctions; Mechanism design; Regret-minimization; Best-response

1 Introduction

We consider problems in the combinatorial auction (CA) domain, where m objects are to be allocated among n potential buyers in order to maximize total value, subject to problem-specific feasibility constraints. These packing problems are complicated by game-theoretic issues: the buyers might benefit from misrepresenting their values to an allocation algorithm. This prompts us to design mechanisms that use payments to encourage reasonable behaviour. The well-known VCG mechanism solves incentive issues by inducing truth-telling as a dominant strategy, but is infeasible for computationally intractible problems. Indeed, for many interesting problems (such as combinatorial auctions), there are large gaps between the best-known approximation factors attainable by efficient truthful mechanisms and those possible in purely computational settings. For some problems, these large gaps are essential [28].

In this extended abstract we consider the problem of designing mechanisms that implement approximation algorithms for combinatorial auction problems without the use of dominant-strategy truthfulness. We are motivated by the domain of repeated auctions, where an auction problem is resolved multiple times with the same objects and bidders¹. These include, for example, auctions for advertising slots [14], bandwidth auctions (such as the FCC spectrum auction), and airline landing rights auctions [10]. In these set-

tings a mechanism for the (one-shot) auction problem corresponds to a repeated game to be played by the agents.

The question of how to model agent behaviour in repeated games has been studied extensively in the economic and algorithmic game theory literature (see Chapters 17-21 of [27] and references therein). Many proposed models assume that agents choose strategies (or distributions thereof) at equilibrium, where no agent has incentive to unilaterally deviate. However, as has been noted elsewhere [4, 16], there are a number of reasons to believe such models are unrealistic: equilibria are computationally hard to find in general, and may not exist without the presence of agents who randomize over strategies for no reason other than to preserve the stability of the system. Even when pure equilibria exist, agents may not necessarily converge to an equilibrium (of the single-round game) or agree on which equilibrium (of the extended-form game) to choose. In light of these concerns, we will focus our attention on two models of agent behaviour that do not make equilibrium assumptions, and have gained recent interest in the algorithmic game theory literature.

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, the average *external regret* of each bidder must tend to 0 as the number of auction rounds increases. These regret-minimizing bidders can be seen as agents that learn how to bid intelligently (relative to any fixed

¹One might alternatively allow preferences and participants to change over time, but sufficiently slowly compared to the rate of auction repetition.

strategy benchmark) from the bidding history of past auction iterations. Note that we require no assumptions about the synchrony or asynchrony of updates; arbitrary sets of agents can update their strategies concurrently. The regret-minimization assumption is motivated by the existence of simple, efficient algorithms that minimize external regret for linear optimization problems [21, 22]. Under this model, 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*, introduced by Blum et al [4].

In the second model, we assume that agents choose strategies that are myopic best-responses to the current strategies of the other agents. Such bidding behaviour is best motivated in settings where agents update their declarations asynchronously. 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 update order. This is closely related to the concept of the *price of (myopic) sinking*, introduced by Goemans et al [16].

Our goal in this area of study is to decouple computational issues from game-theoretic concerns. A full (and admittedly ambitious) solution in our domain would be a black-box conversion of a given approximation algorithm into a mechanism that implements² the same approximation ratio, on average over sufficiently many auction rounds, given our models of bidder behaviour. Our primary research question, partially addressed herein, is to what extent such implementations are possible.

1.1 Our Contribution

We design mechanisms that are based on a particular class of approximation algorithms for combinatorial auction problems: those that are monotone and satisfy the loser-independence property. An algorithm is *monotone* if, whenever a bidder can win some set S by declaring a value of v for it, then he could also win any subset of S with any declared value at least v . This monotonicity condition characterizes truthfulness when bidders are single-minded (meaning

that each agent has value for only a single set), but not for general auction problems [24]. Roughly speaking, an algorithm is *loser-independent* if the outcome for an agent depends only on those agents who would win if he did not participate, and on their declared values for their winnings. This extends a notion of loser-independence for single-parameter problems, introduced by Chekuri and Gamzu [8], to general auction problems. This class of algorithms includes many natural algorithms for well-studied packing problems, including greedy algorithms for CAs [24], convex bundle auctions [1], and unsplittable flow problems [7].

Our first main result is that any monotone loser-independent c -approximate algorithm can be implemented as a mechanism with price of total anarchy at most $c + 1$. That is, if agents minimize their external regret, the average social welfare obtained by our mechanism approaches a $(c + 1)$ approximation to the optimal social welfare as the number of rounds increases. Our mechanism is a black-box reduction from an algorithm for a one-shot auction iteration, and the same mechanism is applied each auction round. The form of our mechanism is very simple: on each round, it applies a simple modification to the bidders' declarations, then runs the approximation algorithm on the modified declarations and charges critical prices (i.e. an agent who wins a set pays the smallest amount he could have declared for that set and won it, given the declarations of the other bidders).

Our implementation does not depend on the specific algorithms used by the agents to minimize their regret; only that their regret vanishes as the number of rounds increases. The rate of convergence to our approximation bound will depend on the rate at which the agents' regret vanishes.

We demonstrate that our mechanism is resilient to the presence of byzantine agents, in the following sense. If each agent either applies regret-minimizing strategies or makes arbitrary declarations (but never declares more than his true value for a set), then the mechanism attains a $(c + 1)$ approximation to the optimal welfare *obtainable by the regret-minimizing bidders*. The no-overbidding assumption is necessary, as otherwise a byzantine agent could bid arbitrarily high and prevent any welfare from being obtained. This assumption is also motivated by viewing byzantine agents as players that do not understand how to participate intelligently in the auction and thus likely to bid conservatively.

We then study the best-response model of bidder behaviour, in which we focus specifically on the combinatorial auction problem. We present a mechanism that implements an $O(s)$ approximation for

²Throughout the paper we use the term “implement” in the economic sense of constructing a mechanism that obtains the desired properties when used by rational agents.

cardinality-restricted combinatorial auctions, where set allocations have size at most s . We then extend this to a mechanism that implements an $O(\sqrt{m})$ approximation for general combinatorial auctions. We attain these approximation ratios with high probability as long as the number of rounds is superlinear in n . We point out that while *truthful* mechanisms with similar approximation ratios are known for *single-minded* combinatorial auctions, our results are significant improvements over what is currently known to be achievable with deterministic truthful algorithms for general CAs.

Returning to the general implementation of monotone loser-independent c -approximate algorithms, we conjecture that the black-box reduction used in the regret-minimization setting also implements an $O(c)$ approximation, on average over sufficiently many rounds, in the model of best-response bidders. We leave the resolution of this conjecture as an open problem.

Our results require a mild game-theoretic assumption, which is that bidders will not apply strategies that are (strictly) dominated by easily-found alternatives. This is precisely the assumption that agents choose only algorithmically undominated strategies, as introduced by Babaioff et al [2]. This assumption is discussed further in Section 3.1. Additionally, the mechanisms we introduce for best-response bidders apply a technique known in implementation theory as *virtual implementation*, where an alternative social choice rule is applied with vanishingly small probability [20]. We view this not as an introduction of randomness into the algorithm being implemented, but rather as the introduction of a trembling-hand consideration into the solution concept that encourages reasonable behaviour when best-response agents must distinguish between otherwise equally beneficial strategies.

1.2 Regret Minimization

We now discuss external regret minimization in further detail. The external regret of a sequence of declarations is the difference between the average utility obtained by an agent (i.e. value of goods received minus payment) and the maximum average utility that the agent could have obtained by making a single fixed declaration each round. An online algorithm for generating declarations is regret-minimizing if its regret vanishes as a function of the number of auction rounds.

How should an agent bid in order to minimize his external regret? A simple and efficient algorithm due to Kalai and Vempala [22] solves linear optimiza-

tion problems with regret that vanishes at a rate of $O(1/\sqrt{T})$. Their algorithm requires access to an exact best-response oracle. Kakade et al [21] show how to use a γ -approximate best response oracle to achieve a γ -approximation to the best fixed declaration in hindsight.

The mechanisms we construct in this paper have the property that the strategy selection problem for each agent reduces to a linear optimization problem over the space of desired object sets. Thus, in settings where each agent has only polynomially many desired sets (e.g. when each agent's type is assumed to be representable by a polynomial number of mutually exclusive bids), it is a simple matter to implement efficient best-response oracles and regret-minimization algorithms. In general, however, agents may have exponentially large strategy spaces³. In such cases an efficient regret-minimization algorithm or best-response oracle would likely have to be tailored to the valuation access model (e.g. oracle queries, succinct representations, etc.) and to the structure of the particular problem and algorithm being implemented. The exploration of such issues for specific problem settings is left as an avenue for future research.

1.3 Related Work

Truthful mechanisms for the combinatorial auction problem have been extensively studied. For general CAs, Hastad's well-known inapproximability result [18] shows that it is hard to approximate the problem to within $\Omega(\sqrt{m})$ assuming $NP \neq ZPP$. The best known deterministic truthful mechanism for CAs with general valuations is a bundling auction that attains an approximation ratio of $O(\frac{m}{\sqrt{\log m}})$ [19]. A randomized $O(\sqrt{m})$ -approximate mechanism that is truthful in expectation was given by Lavi and Swamy [23]. Dobzinski, Nisan and Schapira [13] then gave an $O(\sqrt{m})$ -approximate universally truthful randomized mechanism.

Many variations on the combinatorial auction problem have been considered in the literature. Bartal et al [3] give a truthful $O(Bm^{\frac{1}{B-2}})$ mechanism for multi-unit combinatorial auctions with B copies of each object, for all $B \geq 3$. Dobzinski and Nisan [12] construct a truthful 2-approximate mechanism for multi-unit auctions, and Dobzinski and Dughmi [11] construct a randomized FPTAS for the multi-unit auction problem that is truthful in expectation. Many other problems have truthful mechanisms ([7, 24, 26]) when bidders are restricted to being single-minded. Borodin and Lucier [5] study the limited power of

³More specifically, the corresponding linear optimization problem may have exponential dimension.

certain classes of greedy algorithms for truthfully approximating CA problems.

Implementation at equilibrium, especially for the alternative goal of profit maximization, has a rich history in the economics literature; see, for example, Jackson [20] for a survey. For the goal of optimizing social welfare, Christodoulou et al [9] consider implementing a combinatorial auction by simultaneous Vickrey auctions, and show that this obtains a 2-approximation at every Bayes-Nash equilibrium when agents' valuations are submodular. Gairing et al [15] characterize the Bayes-Nash equilibria of a routing game and study its worst-case performance at equilibrium.

For general combinatorial auction problems, Lucier and Borodin [6] give a black-box reduction from any monotone c -approximate greedy algorithm to a mechanism that obtains a $c+O(\log c)$ approximation at every Bayes-Nash equilibrium. Their results are similar in flavour to our own, though their focus is on single-shot auctions at equilibrium; the process by which such an equilibrium is reached is left open. By contrast, we study the evolution of bidder behaviour in repeated auctions, and demonstrate that it is possible to implement approximation algorithms in settings where convergence to equilibrium is not guaranteed.

The study of regret-minimization goes back to the work of Hannan on repeated two-player games [17]. Kalai and Vempala [22] extend the work of Hannan to online optimization problems, and Kakade et al [21] further extend to settings of approximate regret minimization. Blum et al [4] apply regret-minimization to the study of inefficiency in repeated games, coining the phrase “price of total anarchy” for the worst-case ratio between the optimal objective value and the average objective value when agents minimize regret.

Properties of best-response dynamics in repeated games, and especially the question of convergence to a pure equilibrium, is well-studied (see Chapter 19 of [27]). The study of average performance of best-response dynamics as a metric of game inefficiency, the so-called “price of sinking,” was introduced by Goemans et al [16].

Babaioff et al [2] study implementation of algorithms in undominated strategies, which is a relaxation of the dominant strategy truthfulness concept. They focus on a variant of the CA problem in which agents are assumed to have “single-value” valuations, and present a mechanism to implement such auctions in a multi-round fashion. By comparison, mechanisms in our proposed model solve each instance of an auction in a one-shot manner, and our solution concept assumes that the auction is repeated multiple times.

1.4 Organization

We review mechanism design fundamentals and define relevant problem and algorithm classes in Section 2. In Section 3 we present our general reduction for the regret-minimization model of agent behaviour. We then consider the best-response bidder model in Section 4, where we present mechanisms for the combinatorial auction problem. Conclusions and open problems are discussed in Section 5. Some proofs have been omitted from this extended abstract and may be found in the full version of the paper.

2 Model and Definitions

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, for example, $\mathbf{v} = (v_i, \mathbf{v}_{-i})$.

We consider the domain of combinatorial auction problems, where n agents desire subsets of a set M of m objects. An *allocation profile* is a collection of subsets X_1, \dots, X_n , where X_i is thought of as the subset allocated to agent i . A particular problem instance is defined by the set of *feasible allocation profiles* that are permitted. For example, the general combinatorial auction problem requires that all allocated subsets be disjoint. Each agent i has a privately-held *valuation function* $t_i : 2^M \rightarrow \mathbb{R}$, his *type*, that assigns a value to each allocation. We assume that valuation functions are monotone and normalized so that $v(\emptyset) = 0$. A valuation function v is *single-minded* if there exists $S \subseteq M$ and $x \geq 0$ such that $v(T) = x$ if $S \subseteq T$ and $v(T) = 0$ otherwise. We will write \emptyset for the zero valuation, and (S, x) for a single-minded declaration for S at value x .

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

An allocation rule is *loser-independent* if, whenever $\mathbf{v}_{-i}, \mathbf{v}'_{-i}$ satisfy $\mathcal{A}(\emptyset, \mathbf{v}_{-i}) = \mathcal{A}(\emptyset, \mathbf{v}'_{-i})$ and $v_j(\mathcal{A}_j(\emptyset, \mathbf{v}_{-i})) = v_j(\mathcal{A}_j(\emptyset, \mathbf{v}'_{-i}))$ for all $j \neq i$, then $\mathcal{A}(v_i, \mathbf{v}_{-i}) = \mathcal{A}(v_i, \mathbf{v}'_{-i})$. In other words, agent i 's perception of the behaviour of \mathcal{A} depends only on those agents who would win if agent i did not participate, and on their declared values for their winnings.

A payment rule P assigns a vector of n payments to each valuation profile. A *direct revelation mechanism* \mathcal{M} is composed of an allocation rule \mathcal{A} and a payment rule P . The mechanism proceeds by eliciting a valuation profile \mathbf{d} of *declarations* from the agents, then applying the allocation and payment rules to \mathbf{d} . The utility of agent i for mechanism \mathcal{M} , given declaration

profile \mathbf{d} , is $u_i(\mathbf{d}) = t_i(\mathcal{A}_i(\mathbf{d})) - P_i(\mathbf{d})$. We think of each agent as wanting to choose d_i to maximize $u_i(\mathbf{d})$.

The *social welfare* obtained by allocation profile \mathbf{X} , given type profile \mathbf{t} , is $SW(\mathbf{X}, \mathbf{t}) = \sum_i t_i(X_i)$. Given fixed type profile \mathbf{t} , we write SW_{opt} for $\max_{\mathbf{X}} \{SW(\mathbf{X}, \mathbf{t})\}$, and $SW_{\mathcal{A}}(\mathbf{d}) = \sum_i t_i(\mathcal{A}(\mathbf{d}))$. When $D = (\mathbf{d}^1, \mathbf{d}^2, \dots, \mathbf{d}^T)$ is a sequence of valuation profiles, we write $SW_{\mathcal{A}}(D) = \frac{1}{T} \sum_t SW_{\mathcal{A}}(\mathbf{d}^t)$ for the average welfare obtained over all declarations in D . We will sometimes replace subscript \mathcal{A} by \mathcal{M} , in which case the social welfare is for the allocation rule of \mathcal{M} . Note that algorithm \mathcal{A} is a c -approximation if $SW_{\mathcal{A}}(\mathbf{t}) \geq \frac{1}{c} SW_{opt}$ for all \mathbf{t} .

Given allocation rule \mathcal{A} , agent i , declaration profile \mathbf{d}_{-i} , and set S , the *critical price* $\theta_i^{\mathcal{A}}(S, \mathbf{d}_{-i})$ for S is the minimum value that agent i could bid on set S and be allocated S by \mathcal{A} given fixed \mathbf{d}_{-i} . That is, $\theta_i^{\mathcal{A}}(S, \mathbf{d}_{-i}) = \inf\{v : \exists d_i, d_i(S) = v, \mathcal{A}_i(d_i, \mathbf{d}_{-i}) = S\}$.

We say that a declaration d_i is *weakly dominated* by declaration d_i' for agent i if $u_i(d_i, \mathbf{d}_{-i}) \leq u_i(d_i', \mathbf{d}_{-i})$ for all \mathbf{d}_{-i} , and furthermore there exists some \mathbf{d}'_{-i} such that $u_i(d_i, \mathbf{d}'_{-i}) < u_i(d_i', \mathbf{d}'_{-i})$.

Declaration sequence $D = (\mathbf{d}^0, \mathbf{d}^1, \dots, \mathbf{d}^T)$ *minimizes external regret* for agent i if, for any fixed declaration d_i , $\sum_t u_i(d_i^t, \mathbf{d}_{-i}^t) \geq \sum_t u_i(d_i, \mathbf{d}_{-i}^t) + o(T)$. That is, the utility of agent i approaches the utility of the optimal fixed strategy in hindsight.

Declaration sequence $D = (\mathbf{d}^0, \mathbf{d}^1, \dots, \mathbf{d}^T)$ is an instance of *response dynamics* if, for all $1 \leq t \leq T$, profiles \mathbf{d}^{t-1} and \mathbf{d}^t differ on the declaration of at most one player. Response dynamics D is an instance of *best-response dynamics* if, whenever \mathbf{d}^{t-1} and \mathbf{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, $d_i^t \in \arg \max_d \{u_i(d, \mathbf{d}_{-i}^t)\}$.

3 Regret-Minimizing Bidders

In this section we prove that if agents avoid algorithmically dominated strategies and minimize external regret, then a loser-independent monotone algorithm \mathcal{A} can be converted into a mechanism with almost no loss to its average approximation ratio over sufficiently many rounds. The mechanism, $\mathcal{M}_{\mathcal{A}}$, is described in Figure 1. Mechanism $\mathcal{M}_{\mathcal{A}}$ 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.

3.1 Strategy Selection

The simplification process **SIMPLIFY** essentially converts any declaration into a single-minded declaration (and does not affect declarations that are already

Mechanism $\mathcal{M}_{\mathcal{A}}$:

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

1. $d' \leftarrow \text{SIMPLIFY}(\mathbf{d})$.
2. Allocate $\mathcal{A}(d')$, charge critical prices.

Procedure **SIMPLIFY:**

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

1. For each $i \in [n]$:
2. Choose $S_i \in \arg \max_S \{d_i(S)\}$, breaking ties in favour of smaller sets.
3. $d'_i \leftarrow (S_i, d_i(S_i))$.
4. Return (d'_1, \dots, d'_n) .

Figure 1: Mechanism for regret-minimizing bidders, based on monotone allocation algorithm \mathcal{A} . Uses subprocedure **SIMPLIFY**.

single-minded). We will therefore assume without loss of generality that agents make single-minded declarations, as additional information is not used by the mechanism.⁴

Fix a particular combinatorial auction problem and type profile \mathbf{t} , and let \mathcal{A} be some monotone approximation algorithm. Let \mathbf{d} be a declaration profile; we suppose each d_i is a single-minded bid for set S_i . We draw the following conclusion about the bidding choices of rational agents.

Lemma 3.1. *Declaration d_i is an undominated strategy for agent i if and only if $d_i(S_i) = t_i(S_i)$.*

Proof. For all \mathbf{d}_{-i} , $\mathcal{M}_{\mathcal{A}}(d_i, \mathbf{d}_{-i})$ 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 $t_i(S_i) - \theta_i^{\mathcal{M}_{\mathcal{A}}}(S_i, \mathbf{d}_{-i})$ when $d_i(S_i) > \theta_i^{\mathcal{M}_{\mathcal{A}}}(S_i, \mathbf{d}_{-i})$, and 0 otherwise. A declaration of $d_i(S_i) = t_i(S_i)$ therefore maximizes $u_i(d_i, \mathbf{d}_{-i})$ for all \mathbf{d}_{-i} .

On the other hand, if $d_i(S_i) \neq t_i(S_i)$, let d_i' be the single-minded declaration for S_i at value $t_i(S_i)$. Then for any \mathbf{d}_{-i} such that $\theta_i^{\mathcal{A}}(S_i, \mathbf{d}_{-i})$ lies between $d_i(S_i)$ and $t_i(S_i)$, $u_i(d_i', \mathbf{d}_{-i}) > u_i(d_i, \mathbf{d}_{-i})$. For simplicity we will assume such a \mathbf{d}_{-i} exists; handling the general case requires only a technical and uninteresting extension of notation⁵. Declaration d_i' therefore dominates declaration d_i .

⁴We note, however, that this is not the same as assuming that agents are single-minded; our results hold for bidders with general private valuations.

⁵If $\theta_i^{\mathcal{A}}(S_i, \mathbf{d}_{-i})$ never lies between $d_i(S_i)$ and $t_i(S_i)$ 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

Recall our assumption that agents apply only algorithmically undominated strategies; Lemma 3.1 then implies that each agent will declare her true value for the set on which she makes her single-minded bid. We note that this assumption is not without loss of generality for regret-minimizing bidders, even in mechanisms that are dominant strategy truthful. For example, in a Vickrey auction of a single object, the declaration profile in which one agent bids much higher than anyone's true value and all other agents bid 0 forms an equilibrium, and thus all agents experience zero regret. However, as has been argued elsewhere [6, 9, 25], such strategy profiles seem unnatural: the overbidding agent risks obtaining negative utility if another agent changes declaration, and the underbidding agents gain nothing by bidding 0 instead of their true values. We believe it is reasonable to assume that agents will avoid such risky behaviour and restrict themselves to undominated strategies.

One implication of Lemma 3.1 is that the strategic choice of an agent participating in mechanism $\mathcal{M}_{\mathcal{A}}$ reduces to a linear optimization problem. On each round, we can think of agent i as choosing set S_i , which is the set he will attempt to win that round. Once S_i is chosen, an undominated declaration for agent i is determined: the single-minded declaration for S_i at value $t_i(S_i)$. Given that agent i chooses set S_i , his utility will be $t_i(S_i) - w_i$, where $w_i = \min\{t_i(S_i), \theta_i^{\mathcal{A}}(S_i, \mathbf{d}_{-i})\}$ is the price for set S_i , determined by the declarations of the other agents, capped at $t_i(S_i)$. As a corollary, if an agent's valuation is representable as a polynomial number of mutually exclusive bids, then the regret-minimization algorithm of Kalai and Vempala [22] can be used to efficiently choose strategies that minimize external regret.

3.2 Performance of $\mathcal{M}_{\mathcal{A}}$

We now proceed with bounding the social welfare obtained by $\mathcal{M}_{\mathcal{A}}$. Let A_1, \dots, A_n be an optimal assignment for types \mathbf{t} . Suppose that $D = \mathbf{d}^1, \dots, \mathbf{d}^T$ is a sequence of declarations to our mechanism. The definition of regret minimization then immediately implies the following.

Lemma 3.2. *If agent i minimizes her external regret in bid sequence D , then $\frac{1}{T} \sum_t (t_i(\mathcal{A}(\mathbf{d}^t)) + \theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}^t)) \geq t_i(A_i) - o(1)$.*

Proof. Let d_i^t be the single-minded declaration for set A_i at value $t_i(A_i)$. From the definition of regret min-

being "the same" as a single-minded declaration for S_i at value $t_i(S_i)$. We will ignore this technical issue for the remainder of the paper, in the interest of keeping the exposition simple.

imization,

$$\begin{aligned} \frac{1}{T} \sum_t u_i(d_i^t, \mathbf{d}_{-i}^t) &\geq \frac{1}{T} \sum_t u_i(d_i^t, \mathbf{d}_{-i}^t) - o(1) \\ &\geq \frac{1}{T} \sum_t (t_i(A_i) - \theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}^t)) - o(1) \\ &= t_i(A_i) - \frac{1}{T} \sum_t \theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}^t) - o(1). \end{aligned}$$

Since $u_i(d_i^t, \mathbf{d}_{-i}^t) \leq t_i(\mathcal{A}(\mathbf{d}^t))$ for all t , the result follows.

Assume now that algorithm \mathcal{A} is loser independent. We can then relate the value of the solution returned by an algorithm to the critical prices of the sets in an optimal solution.

Lemma 3.3. *If \mathcal{A} is a monotone loser-independent c -approximate algorithm, then $\sum_i d_i(\mathcal{A}(\mathbf{d})) \geq \frac{1}{c} \sum_i \theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i})$.*

Proof. Choose $\epsilon > 0$. For each i , let d_i' be the pointwise maximum between d_i and the single-minded declaration for set A_i at value $\theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}) - \epsilon$. The definition of loser independence implies that critical prices are the same under declaration profiles \mathbf{d} and \mathbf{d}' , and moreover $\mathcal{A}(\mathbf{d}') = \mathcal{A}(\mathbf{d})$. Since \mathcal{A} is a c -approximate algorithm, $\sum_i d_i'(\mathcal{A}(\mathbf{d}')) \geq \frac{1}{c} \sum_i d_i'(A_i) \geq \frac{1}{c} \sum_i (\theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}) - \epsilon)$. Additionally, since $d_i'(T) = d_i(T)$ whenever $d_i(T) \geq \theta_i^{\mathcal{A}}(T, \mathbf{d}_{-i})$ (from the definition of d_i'), we have $d_i'(\mathcal{A}(\mathbf{d}')) = d_i(\mathcal{A}(\mathbf{d}))$ for all i . We conclude that $\sum_i d_i(\mathcal{A}(\mathbf{d})) \geq \frac{1}{c} \sum_i (\theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}) - \epsilon)$ for all $\epsilon > 0$, as required.

We are now ready to proceed with the proof of our main result in this section.

Theorem 3.4. *Any monotone loser-independent c -approximate algorithm can be implemented as a mechanism with $c + 1$ price of total anarchy.*

Proof. Let $D = \mathbf{d}^1, \dots, \mathbf{d}^T$ be a sequence of declarations in which all agents minimize external regret. By Lemma 3.2, $\frac{1}{T} \sum_t (t_i(\mathcal{A}(\mathbf{d}^t)) + \theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}^t)) \geq t_i(A_i) - o(1)$. Summing over all i , we have $\frac{1}{T} \sum_t \sum_i (t_i(\mathcal{A}(\mathbf{d}^t)) + \theta_i^{\mathcal{A}}(A_i, \mathbf{d}_{-i}^t)) \geq SW_{OPT} - (n)(o(1))$. By Lemma 3.3, this implies $\frac{1}{T} \sum_t \sum_i (t_i(\mathcal{A}(\mathbf{d}^t)) + cd_i^t(\mathcal{A}(\mathbf{d}^t))) \geq SW_{OPT} - (n)(o(1))$. We know $d_i^t(\mathcal{A}(\mathbf{d}^t)) = t_i(\mathcal{A}(\mathbf{d}^t))$ for all i and t by Lemma 3.1, so we conclude

$$\begin{aligned} (c + 1)SW_{\mathcal{A}}(D) &= (c + 1) \frac{1}{T} \sum_t \sum_i t_i(\mathcal{A}(\mathbf{d}^t)) \\ &\geq SW_{OPT} - (n)(o(1)). \end{aligned}$$

Since the term hidden by the asymptotic notation vanishes with T and does not depend on n , we obtain the desired result.

Theorem 3.4 is very general, as it applies to a number of known algorithms for various problem settings. For example, Theorem 3.4 yields an $O(\sqrt{m})$ implementation of the combinatorial auction problem [24], an $s + 1$ implementation of the combinatorial auction problem where sets are restricted to cardinality s (using a simple greedy algorithm), an $O(Bm^{1/(B-1)})$ implementation of the unsplittable flow problem with minimum edge capacity B [7], and an $O(R^{4/3})$ implementation of the combinatorial auction of convex bundles in the plane where R is the maximum aspect ratio over all desired bundles [1].

We note that, since agents experience no regret at a pure Nash equilibrium, an immediate corollary to Theorem 3.4 is that any monotone loser-independent c -approximate algorithm can be implemented as a mechanism with $c + 1$ price of anarchy. We remark that an alternative proof of this result has been given recently using a different mechanism construction [6].

Also, the rate at which the welfare obtained by \mathcal{M}_A converges to an average that is a $c + 1$ approximation to optimal depends on the rate of convergence of players' external regret to 0. The average welfare obtained after T rounds will have an additive loss of $(n)r(T)$, where $r(T)$ is the average regret experienced by an agent after T rounds. Assuming that agents apply algorithms that minimize regret at a rate of $r(T) = o(1/\sqrt{T})$, which is attainable using the algorithm of Kalai and Vempala [22], the additive error term is at most a constant when T is at least quadratic in n .

3.3 Resilience to Byzantine Agents

Suppose that in addition to regret-minimizing agents, the auction participants include byzantine agents. The only restriction we impose on the behaviour of such agents is that they do not overbid on any set; that is, $d_i(S) \leq t_i(S)$ for any S and byzantine agent i . We can motivate this restriction either through our characterization of undominated strategies in Lemma 3.1, or by thinking of byzantine players as not understanding how to participate rationally in the auction, and hence likely to be conservative in the way that they bid. Under this assumption, since Lemma 3.3 holds for any declaration profile, we easily obtain the following generalization of Theorem 3.4.

Proposition 3.5. *Suppose \mathcal{A} is a monotone loser-independent c -approximate algorithm and D is a declaration sequence for \mathcal{M}_A . If $N \subseteq [n]$ is a collection of agents that minimize regret in D , and*

the remaining agents never bid more than their true values on any set in D , then $\frac{1}{T} \sum_t SW_A(\mathbf{d}^t) \geq \frac{1}{c+1} \sum_{i \in N} SW_{opt} + |N|o(1)$.

3.4 Importance of Loser-Independence

We note that the loser independence property is necessary for Theorem 3.4, as the following example demonstrates.

Example 3.6. Consider an auction problem in which no agent can be allocated more than s objects, and moreover $M = A \cup B$ where $|A| = |B| = m/2$ and the mechanism can either allocate objects in A or objects in B , but not both. Consider the algorithm that takes the maximum over two solutions: a greedy assignment of subsets of A , and a greedy assignment of subsets of B . This algorithm obtains an $s + 1$ approximation, but is not loser-independent.

Consider now an instance of the problem in which a single agent desires all of B with value 1, and each of $m/2$ agents desires a separate singleton in A with value $1 - \epsilon$. Suppose that the agent desiring B declares his valuation truthfully, but the other agents declare the zero valuation. On this input, the algorithm under consideration obtains only an $m/2$ approximation to the optimal solution. However, this set of declarations forms a Nash equilibrium, and hence each agent has zero regret under this input profile. Thus, even if agents minimize their regret, our mechanism may obtain a very poor approximation to the optimal social welfare over arbitrarily many auction rounds.

4 Best-Response Agents

In this section we consider the problem of designing mechanisms for agents that apply myopic best-response strategies asynchronously. Recall that in our model agents are chosen for update uniformly at random, one per round. In order to keep our exposition clear, we will make two additional assumptions about the nature of the best-response behaviour (which can be removed, as we discuss in Section 4.3). First, we will suppose that in the initial state every bidder makes the empty declaration \emptyset . Second, we suppose that if a bidder is chosen for update but cannot improve his utility, he will choose to maintain his previous strategy. These assumptions will simplify the process of characterizing best-response strategies of agents, and in particular the statement of Lemma 4.4 in the next section. It is possible to remove these assumptions, at the cost of a minor modification to the mechanisms we propose. We defer a more complete discussion to Section 4.3.

We begin our analysis of myopic bidding strategies by considering mechanism \mathcal{M}_A from Section 3 (given

a monotone loser-independent algorithm \mathcal{A}). It is tempting to guess that the best-response dynamics for this mechanism will necessarily converge to equilibrium, but the following example shows that this is not the case.

Example 4.1. 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 $\mathcal{M}_{\mathcal{A}}$, and agents apply best-response dynamics. Agents 3 and 4 are single-minded and hence 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 each agent bids is determined by Lemma 3.1 (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, starting from one of these four declaration profiles⁶, the best-response dynamics will never converge to equilibrium.

The example above motivates a study of the *average* social welfare of $\mathcal{M}_{\mathcal{A}}$ over multiple rounds of best-response dynamics. We conjecture that, on average, the best-response dynamics on mechanism $\mathcal{M}_{\mathcal{A}}$ approximates the optimal social welfare to within a constant factor of the approximation ratio of the original algorithm \mathcal{A} .

Conjecture 4.2. *If \mathcal{A} is a monotone loser-independent c -approximate algorithm, then $\mathcal{M}_{\mathcal{A}}$ has $O(c)$ price of (myopic) sinking.*

We leave the resolution of Conjecture 4.2 as an open problem. As partial progress, we construct alternative mechanisms that are more amenable to best-response analysis. These mechanisms are tailored specifically

⁶Such a profile can be reached from the initial empty state by choosing agent 3 for update, followed by 4, 1, and then 2.

to the general combinatorial auction problem and to combinatorial auctions with cardinality-restricted sets.

The primary tool we will use is the following probabilistic lemma, which pertains to any mechanism in a best-response setting. Suppose \mathcal{M} is a mechanism, and D is a sequence of best-response declarations for \mathcal{M} . For any \mathbf{d} , let $P_1(\mathbf{d}) = P_1(\mathbf{d}_{-i})$ be some property of \mathbf{d} that does not depend on d_i , and let $P_2(\mathbf{d}) = P_2(d_i)$ be some property depending only on d_i .

Lemma 4.3. *Suppose that, for any \mathbf{d} , if $P_1(\mathbf{d}_{-i})$ is false, then any best response by agent i , d_i , satisfies $P_2(d_i)$. Then for all $\epsilon > 0$, if best-response dynamics is run for $T > \epsilon^{-1}n$ steps, there will be at least $(\frac{1}{2} - \epsilon)T$ steps t for which either $P_1(\mathbf{d}_{-i}^t)$ or $P_2(d_i^t)$ is true, with probability at least $1 - e^{-T\epsilon^2/32n}$.*

Proof (sketch). Let B_i^t be the event that neither $P_1(\mathbf{d}_{-i}^t)$ nor $P_2(d_i^t)$ is true, and let A_i^t denote the event that $P_2(d_i^t)$ is true. Our goal is to bound the number of occurrences of B_i^t .

Note that if B_i^t occurs and agent i is chosen for update on step $t+1$, then A_i^{t+1} occurs (by assumption). Alternatively, if A_i^t occurs but agent i is not chosen for update on step $t+1$ then A_i^{t+1} occurs, since A_i depends only on the declaration of agent i . Events A_i^t and B_i^t can therefore be compared to a random walk on $\{0, 1\}$, where at each step the current state changes with probability $1/n$. The number of occurrences of B_i^t is dominated by the number of occurrences of 0 in such a random walk. A straightforward application of the method of bounded average differences shows that this value is concentrated around its expectation, which is at most $\frac{T}{2} + \frac{n}{2}$. Thus, as long as $T > \epsilon^{-1}n$, the number of occurrences of B_i^t will be at most $T(\frac{1}{2} + \epsilon)$ with high probability, giving the desired bound. Additional details appear in the full version of the paper.

4.1 A Mechanism for s -CAs

Consider the s -CA problem, which is a combinatorial auction in which no agent can be allocated more than s objects. An algorithm that greedily assigns sets in descending order by value obtains an $(s+1)$ approximation.⁷ Call this algorithm \mathcal{A}_{sCA} . We will construct a mechanism \mathcal{M}_{sCA} based on \mathcal{A}_{sCA} ; it is described in Figure 2. This algorithm simplifies incoming bids (in the same way as $\mathcal{M}_{\mathcal{A}}$) 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 larger

⁷And an s approximation for single-minded declarations.

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. $R_i \leftarrow \{j : S_j \cap T_i \neq \emptyset\}$.
5. $p_i \leftarrow \sum_{j \in R_i} d_j(S_j)$.
6. If $d_i'(T_i) \leq p_i$, set $T_i \leftarrow \emptyset$, $p_i \leftarrow 0$.
7. Allocate T_1, \dots, T_n , charge critical prices.

Figure 2: Mechanism \mathcal{M}_{sCA} , an implementation of greedy algorithm \mathcal{A}_{sCA} for the s -CA problem.

than 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 claim that \mathcal{M}_{sCA} obtains an $O(s)$ approximation to the optimal social welfare on average over sufficiently many rounds of best-response dynamics.

We note that since our mechanism implements a monotone algorithm and charges critical prices, Lemma 3.1 implies that undominated strategies for agent i involve choosing a set S_i and making a single-minded bid for S_i at value $t_i(S_i)$. We will therefore assume that agents bid in this way.

Suppose that \mathbf{d} is a declaration profile, where each d_i is a single-minded declaration for some set S_i . For any set $T \subseteq M$, define $R_i(\mathbf{d}, T) = \{j : j \neq i, S_j \cap T \neq \emptyset\}$. We also define $Q_i(\mathbf{d}, T) = \{j : j \in R_i(\mathbf{d}, T), d_j(S_j) < t_i(T)\}$. That is, $R_i(\mathbf{d}, T)$ is the set of bidders other than i whose single-minded declared sets intersect T , and $Q_i(\mathbf{d}, T)$ is the subset of those bidders whose single-minded declared values are less than agent i 's true value for T . Note that R_i on line 4 of \mathcal{M}_{sCA} is precisely $R_i(\mathbf{d}', T_i)$. We say that \mathbf{d} is *separated for agent i* if $\sum_{j \in Q_i(\mathbf{d}, S_i)} d_j(S_j) \leq d_i(S_i)$ and \mathbf{d} is *separated* if it is separated for every bidder. Since an agent gains positive utility only if the declaration is separated for him, and since the initial (empty) declaration profile is separated, we draw the following conclusion.

Lemma 4.4. *At each step of the best-response dynamics for mechanism \mathcal{M}_{sCA} , the declaration profile submitted by the agents will be separated.*

For the remainder of the section we will assume that declaration profiles are separated. Under this assumption,

the behaviour of mechanism \mathcal{M}_{sCA} simplifies in a fortuitous way.

Proposition 4.5. *If \mathbf{d} is separated, then \mathcal{M}_{sCA} allocates S_i to agent i precisely when $d_i(S_i) > \max_{j \in R_i(S_i, \mathbf{d})} d_j(S_j)$.*

Proof. Note that \mathcal{M}_{sCA} allocates S_i to agent i if and only if \mathcal{A}_{sCA} allocates S_i to agent i on input \mathbf{d}' and $d_i(S_i) > \sum_{j \in R_i(\mathbf{d}, S_i)} d_j(S_j)$.

Suppose that $d_i(S_i) > \max_{j \in R_i(S_i, \mathbf{d})} d_j(S_j)$. Then $\mathcal{A}_{sCA}(\mathbf{d}')$ allocates S_i to agent i , and furthermore $Q_i(S_i, \mathbf{d}) = R_i(S_i, \mathbf{d})$. Since \mathbf{d} is separated, $d_i(S_i) > \sum_{j \in Q_i(S_i, \mathbf{d})} d_j(S_j)$ and therefore $d_i(S_i) > \sum_{j \in R_i(S_i, \mathbf{d})} d_j(S_j)$. We conclude that S_i will be allocated to agent i by \mathcal{M}_{sCA} . On the other hand, if $d_i(S_i) \leq \max_{j \in R_i(S_i, \mathbf{d})} d_j(S_j)$, then certainly $d_i(S_i) \leq \sum_{j \in R_i(S_i, \mathbf{d})} d_j(S_j)$, so S_i is not allocated to agent i by \mathcal{M}_{sCA} .

Let A_1, \dots, A_n be an optimal allocation with respect to the agents' true types \mathbf{t} .

Proposition 4.6. *If \mathbf{d} is separated and $\sum_{j \in R_i(\mathbf{d}, A_i)} d_j(S_j) < \frac{1}{2}t_i(A_i)$, then any utility-maximizing declaration for agent i , d_i , will be a single-minded declaration for some S_i with $d_i(S_i) \geq \frac{1}{2}t_i(A_i)$.*

Proof. It can be verified that $\theta_i^{\mathcal{M}_{sCA}}(A_i, \mathbf{d}_{-i}) = \sum_{j \in R_i(\mathbf{d}, A_i)} d_j(S_j)$, so agent i would obtain utility at least $\frac{1}{2}t_i(A_i)$ by making a single-minded declaration for set A_i at value $t_i(A_i)$. His utility-maximizing declaration must therefore make at least this much utility, and hence is a bid for some set S_i with $d_i(S_i) = t_i(S_i) \geq \frac{1}{2}t_i(A_i)$.

For declaration profile \mathbf{d} , let G denote the set of agents i for which either $\sum_{j \in R_i(\mathbf{d}, A_i)} d_j(S_j) \geq \frac{1}{2}t_i(A_i)$ or $d_i(S_i) \geq \frac{1}{2}t_i(A_i)$. We can then bound the social welfare obtained by \mathcal{M}_{sCA} with respect to the optimal assignment to agents in G .

Lemma 4.7.

$$SW_{\mathcal{M}_{sCA}}(\mathbf{d}) \geq \frac{1}{4(s+1)} \sum_{i \in G} t_i(A_i).$$

We are now ready to bound the average social welfare of our mechanism, over sufficiently many rounds, with respect to the approximation factor of algorithm \mathcal{A} .

Theorem 4.8. *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

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

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

Proof. Let G_t be the set of agents G from Lemma 4.7 on step t (i.e. with respect to declaration \mathbf{d}^t). Lemma 4.3 and Proposition 4.6 together imply 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 4.7 implies

$$\begin{aligned} SW_{\mathcal{M}_{sCA}}(D) &= \frac{1}{T} \sum_t SW_{\mathcal{M}_{sCA}}(\mathbf{d}^t) \\ &\geq \frac{1}{4(s+1)T} \sum_t \sum_{i \in G_t} t_i(A_i) \\ &\geq \frac{1}{4(s+1)T} \sum_i T \left(\frac{1}{2} - \epsilon \right) t_i(A_i) \\ &\geq \left(\frac{1-2\epsilon}{8(s+1)} \right) SW_{opt}(\mathbf{t}). \end{aligned}$$

as required.

If we take ϵ to be a small constant and assume $T = \Omega(n^{1+\delta})$ for some $\delta > 0$, we conclude that $SW_{\mathcal{M}_{sCA}}(D) > \frac{1}{O(s)} SW_{opt}(\mathbf{t})$ with high probability. Thus \mathcal{M}_{sCA} implements an $O(s)$ approximation to the s -CA problem for best-response bidders, on average, when the number of rounds is superlinear in n .

4.2 A Mechanism for General CAs

Consider the following algorithm for the general CA problem: try greedily assigning sets, of size at most \sqrt{m} , by value; return either the resulting solution or the allocation that gives all items to a single agent, whichever generates more welfare. This algorithm is an $O(\sqrt{m})$ approximation [26]. We will construct a mechanism \mathcal{M}_{CA} based on this algorithm; it is described in Figure 3. \mathcal{M}_{CA} essentially implements two copies of \mathcal{M}_{sCA} : 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. \mathcal{M}_{CA} then takes the maximum of the two solutions. We add one additional modification: with vanishingly small probability γ , \mathcal{M}_{CA} ignores bids for M and behaves as $\mathcal{M}_{\sqrt{m}CA}$. The purpose of this modification is to encourage agents to bid on small sets, even when the

Mechanism \mathcal{M}_{CA} :

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. With probability γ :
3. For all i with $S_i = M$, $d'_i \leftarrow \emptyset$.
4. Let $(T_1, \dots, T_n) \leftarrow \mathcal{M}_{\sqrt{m}CA}(\mathbf{d}')$.
5. If $\exists i : S_i = M$:
6. Let $j \leftarrow \arg \max_j \{d'_j(M) : S_j = M\}$.
7. If $d'_j(S_j) > \sum_{k \neq j: S_k = M} d'_k(S_k)$ and $d'_j(S_j) > \sum_i d'_i(T_i)$:
8. Set $T_j \leftarrow M$, $T_i \leftarrow \emptyset$ for all $i \neq j$
9. Allocate T_1, \dots, T_n , charge critical prices.

Figure 3: Mechanism \mathcal{M}_{CA} , a best-response implementation of a greedy algorithm for the CA problem. Parameter $\gamma > 0$ is an arbitrarily small positive constant. Note $\mathcal{M}_{\sqrt{m}CA}$ is \mathcal{M}_{sCA} from Figure 2 with $s = \sqrt{m}$.

presence of a high-valued bid for a large set would seem to indicate that bidding on small sets is fruitless.

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 primary complication is that the bidding choice of an agent may be influenced by the mechanism's choice of whether or not to allocate M to a single bidder; this can be handled by a careful analysis of utility-maximizing declarations. The details are deferred to the full version of the paper. We obtain the following result.

Theorem 4.9. *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*

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

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

If we take ϵ to be a small constant and assume $T = \Omega(n^{1+\delta})$ for some $\delta > 0$, we conclude that $SW_{\mathcal{M}_{CA}}(D) > \frac{1}{O(\sqrt{m})} SW_{opt}(\mathbf{t})$ with high probability. Thus \mathcal{M}_{CA} implements an $O(\sqrt{m})$ approximation to the general CA problem for best-response bidders, on average, when the number of rounds is superlinear in n .

4.3 Removing Additional Assumptions

Recall that we made two additional assumptions in our model of best-response dynamics: that each bidder makes the empty declaration \emptyset in the initial state, and that if a bidder is chosen for update, but cannot improve his utility, he will choose to maintain his previous strategy. We used these assumptions to argue that agents make only separated declarations when participating in mechanisms \mathcal{M}_{sCA} and \mathcal{M}_{CA} .

These assumptions can be removed as follows. We modify mechanisms \mathcal{M}_{sCA} and \mathcal{M}_{CA} so that, with vanishingly small probability, an alternative allocation rule is used⁸. This alternative rule chooses an agent at random and assigns him all objects (or a randomly chosen maximal feasible set of objects) at no cost *as long as the input declaration is separated for that agent*. Under this modified mechanism, any separated declaration by agent i results in positive expected utility. Since any non-separated declaration results in a utility of 0, it must be that the utility-maximizing declaration by any agent must be separated.

It follows that after each bidder is chosen at least once for update, and every step thereafter, the input declaration will be separated. Thus, with high probability, every declaration after $O(n \log n)$ steps will be separated. Lemma 4.4 will therefore hold after $O(n \log n)$ steps of best-response dynamics, with high probability; the remainder of the analysis then proceeds as before.

5 Conclusions and Future Work

We considered the problem of designing mechanisms for use with regret-minimizing and best-response bidders in repeated combinatorial auctions. We presented a general black-box construction for the regret-minimization model, which implements any monotone loser-independent approximation algorithm. For the best-response model, we constructed an $O(\sqrt{m})$ -approximate mechanism for the combinatorial auction problem.

One direction for future research is to extend our results to implement additional algorithms. Our best-response mechanisms made specific use of the structure of greedy CA algorithms, but it seems likely that our approach can be generalized. It would also be worthwhile to explore efficient algorithms for minimizing regret and determining best-responses when agent

types are exponential (e.g. general combinatorial auctions), possibly tailored to particular problem settings and algorithms. Another question of note is whether Conjecture 4.2 is true, and the mechanism we proposed for regret-minimizing bidders also yields good performance for best-response bidders. More generally, can our techniques be modified to apply to algorithms that are not loser-independent? A broader research topic is to explore other models for reasonable bidder behaviour, which may admit different mechanism implementations.

Acknowledgements

Many thanks go to Allan Borodin and Silvio Micali for helpful discussions, and to the anonymous referees whose comments greatly improved the quality of this paper.

References

- [1] M. Babaioff and L. Blumrosen. Computationally-feasible truthful auctions for convex bundles. In *Proc. 7th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2004.
- [2] M. Babaioff, R. Lavi, and E. Pavlov. Single-value combinatorial auctions and algorithmic implementation in undominated strategies. *Journal of the ACM*, 56(1):1–32, 2009.
- [3] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *Proc. 9th Conf. on Theoretical Aspects of Rationality and Knowledge*, 2003.
- [4] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *Proc. 39th ACM Symp. on Theory of Computing*, 2008.
- [5] A. Borodin and B. Lucier. Greedy mechanism design for truthful combinatorial auctions. Working Paper, 2009.
- [6] A. Borodin and B. Lucier. Price of anarchy for greedy auctions. In *Proc. 21th ACM Symp. on Discrete Algorithms*, 2010.
- [7] P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanism design. In *Proc. 36th ACM Symp. on Theory of Computing*, 2005.
- [8] C. Chekuri and I. Gamzu. Truthful mechanisms via greedy iterative packing. In *Proc. 12th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2009.
- [9] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *Proc. 35th*

⁸As discussed in Section 1.1, this perturbation can be viewed as a trembling-hand variation of the best-response solution concept, rather than an introduction of randomness into the allocation algorithms.

- Intl. Colloq. on Automata, Languages and Programming*, pages 820–832, 2008.
- [10] P. Cramton, Y. Shoham, and R. S. (eds.). *Combinatorial Auctions*. MIT Press, Cambridge, MA, 2005.
- [11] S. Dobzinski and S. Dughmi. On the power of randomization in algorithmic mechanism design. *CoRR*, abs/0904.4193, 2009.
- [12] S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *Proc. 9th ACM Conf. on Electronic Commerce*, 2007.
- [13] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proc. 37th ACM Symp. on Theory of Computing*, 2006.
- [14] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In *Stanford Graduate School of Business Research Paper No. 1917*, 2005.
- [15] M. Gairing, B. Monien, and K. Tiemann. Selfish routing with incomplete information. In *Proc. 17th ACM Symp. on Parallel Algorithms and Architectures*, 2005.
- [16] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, 2005.
- [17] J. Hannan. Approximation to bayes risk in repeated plays. In M. Dresher, A. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume 4. Princeton University Press, 1957.
- [18] J. Hastad. Some optimal inapproximability results. In *Proc. 29th ACM Symp. on Theory of Computing*, 1997.
- [19] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.
- [20] M. Jackson. A crash course in implementation theory. In *Social Choice and Welfare*, 2001.
- [21] S. Kakade, A. Kalai, and K. Ligett. Playing games with approximation algorithms. In *Proc. 38th ACM Symp. on Theory of Computing*, 2007.
- [22] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 2005.
- [23] R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, 2005.
- [24] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce*, pages 96–102. ACM Press, 1999.
- [25] R. P. Leme and E. Tardos. Sponsored search equilibria for conservative bidders. In *Fifth Workshop on Ad Auctions*, 2009.
- [26] A. Mu’alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64:612–631, 2008.
- [27] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [28] C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, 2008.