

Limitations of Randomized Mechanisms for Combinatorial Auctions

Shaddin Dughmi *
Stanford University
Stanford, CA, USA
shaddin@cs.stanford.edu

Jan Vondrák
IBM Almaden Research Center
San Jose, CA, USA
jvondrak@us.ibm.com

Abstract— The design of computationally efficient and incentive compatible mechanisms that solve or approximate fundamental resource allocation problems is the main goal of algorithmic mechanism design. A central example in both theory and practice is welfare-maximization in combinatorial auctions. Recently, a randomized mechanism has been discovered for combinatorial auctions that is truthful in expectation and guarantees a $(1 - 1/e)$ -approximation to the optimal social welfare when players have coverage valuations [11]. This approximation ratio is the best possible even for non-truthful algorithms, assuming $P \neq NP$ [16].

Given the recent sequence of negative results for combinatorial auctions under more restrictive notions of incentive compatibility [7], [2], [9], this development raises a natural question: Are truthful-in-expectation mechanisms compatible with polynomial-time approximation in a way that deterministic or universally truthful mechanisms are not? In particular, can polynomial-time truthful-in-expectation mechanisms guarantee a near-optimal approximation ratio for more general variants of combinatorial auctions?

We prove that this is not the case. Specifically, the result of [11] cannot be extended to combinatorial auctions with submodular valuations in the value oracle model. (Absent strategic considerations, a $(1 - 1/e)$ -approximation is still achievable in this setting [25].) More precisely, we prove that there is a constant $\gamma > 0$ such that there is no randomized mechanism that is truthful-in-expectation — or even approximately truthful-in-expectation — and guarantees an $m^{-\gamma}$ -approximation to the optimal social welfare for combinatorial auctions with submodular valuations in the value oracle model.

We also prove an analogous result for the flexible combinatorial public projects (CPP) problem, where a truthful-in-expectation $(1 - 1/e)$ -approximation for coverage valuations has been recently developed [13]. We show that there is no truthful-in-expectation — or even approximately truthful-in-expectation — mechanism that achieves an $m^{-\gamma}$ -approximation to the optimal social welfare for combinatorial public projects with submodular valuations in the value oracle model. Both our results present an unexpected separation between coverage functions and submodular functions, which does not occur for these problems without strategic considerations.

1. INTRODUCTION

The design of incentive-compatible mechanisms for welfare maximization in combinatorial auctions is a central problem of algorithmic mechanism design. In a combinatorial auction, there are n players and a set M of m items. Player i has a (private) valuation function $v_i : 2^M \rightarrow \mathbb{R}_+$ which is assumed to be monotone ($v_i(S) \leq v_i(T)$ whenever $S \subset T$) and normalized ($v_i(\emptyset) = 0$). The goal is to

design a computationally efficient mechanism that yields an allocation of items (S_1, \dots, S_n) to the players along with payments (p_1, \dots, p_n) so that (a) the *social welfare* $\sum_{i=1}^n v_i(S_i)$ is approximately maximized, and (b) the mechanism is *incentive-compatible*, or *truthful*, meaning that each player maximizes his *utility* $v_i(S_i) - p_i$ by reporting his true valuation v_i .

This problem has been studied extensively in both strategic and non-strategic settings. Various strategic solution concepts have been considered, including deterministic truthfulness, universal truthfulness, and truthfulness in expectation. Moreover, both strategic and non-strategic formulations of the problem have been studied for various restricted classes of valuations, as well as under various assumptions on how valuations are accessed or represented. Absent assumptions on the class of valuations, the welfare maximization problem is very hard to approximate even by non-truthful algorithms (NP-hardness of $m^{\epsilon-1/2}$ -approximation follows from the set packing problem). Better approximation ratios are possible for valuation classes that restrict complementarity between items. A prominent class of such valuations are *submodular functions*: functions v_i where the marginal value $v_i(S \cup \{j\}) - v_i(S)$ for a each fixed item j is non-increasing in S . It is known that the welfare maximization problem with submodular valuation functions admits a (non-truthful) $(1 - 1/e)$ -approximation algorithm [25], and this is optimal assuming $P \neq NP$ [16]. The hardness result of [16] holds even in the special case of coverage valuations; the algorithmic result of [25] holds in the *value oracle model*, where each v_i can be queried only through an oracle returning $v_i(S)$ for a given query S . This is also the model we consider in this paper. It is known that a $(1 - 1/e + \epsilon)$ -approximation for combinatorial auctions with submodular bidders would require exponentially many value queries [20].

The classical VCG mechanism is incentive compatible and maximizes welfare in combinatorial auctions. Unfortunately, however, VCG can not be implemented in polynomial time even for very special classes of valuation functions, including submodular functions. Combining computational efficiency and truthfulness for combinatorial auctions appears difficult. A series of works have provided evidence that computational efficiency and truthfulness are in conflict: (deterministic) VCG-type mechanisms have been ruled out

*Supported by NSF Grant CCF-0448664.

Class of valuations	Approximation	Universally truthful	Truthful-in-expectation
submodular / value oracle	$1 - 1/e$	$m^{-1/2} \mid m^{\epsilon-1/2}$	$m^{-1/2} \mid m^{-\gamma}$ [new]
coverage, matroid rank sums	$1 - 1/e$	$m^{-1/2} \mid 1 - 1/e$	$1 - 1/e$
budget-additive	$\frac{3}{4} \mid \frac{15}{16}$	$\Omega(1/\log m \log \log m) \mid \frac{15}{16}$	$\Omega(1/\log m \log \log m) \mid \frac{15}{16}$
submodular / demand oracle	$1 - 1/e + \epsilon \mid \frac{15}{16}$	$\Omega(1/\log m \log \log m) \mid \frac{15}{16}$	$\Omega(1/\log m \log \log m) \mid \frac{15}{16}$

Figure 1. Currently known results for combinatorial auctions: approximation \mid inapproximability. If only one result is given, it is known to be optimal. For randomized maximal-in-range (universally truthful) mechanisms, it is known that it is hard to achieve a better than $1/n$ -approximation for coverage valuations; however, other universally truthful mechanisms might exist. No non-trivial hardness was previously known for truthful-in-expectation combinatorial auctions, even when restricted to maximal-in-distributional-range mechanisms.

for submodular combinatorial auctions in the communication complexity model [7], and even for explicitly given budget-additive valuations [2]. Recently, Dobzinski [9] proved that there is no deterministic truthful or even randomized universally truthful mechanism for submodular combinatorial auctions in the value oracle model, achieving an approximation ratio better than $m^{\epsilon-1/2}$.

Therefore, it came as a surprise when a $(1 - 1/e)$ -approximate randomized mechanism was discovered by Dughmi, Roughgarden and Yan [11] for a large subclass of submodular valuations. Their mechanism is *truthful in expectation* — a weaker notion than truthfulness in the universal sense — and applies to explicitly represented coverage functions. More generally, their mechanism applies to “black-box” valuations expressible as weighted sums of matroid rank functions, provided they support “lottery-value queries” (what is the expectation $\mathbf{E}[v_i(\widehat{\mathbf{x}})]$ for a given product distribution $\widehat{\mathbf{x}}$). The mechanism can be also implemented in the value oracle model, at the cost of relaxing the solution concept to approximate truthfulness in expectation [12].

This development raises a natural question: Could truthfulness-in-expectation be the cure for combinatorial auctions, perhaps providing an optimal $(1 - 1/e)$ -approximation for all submodular valuations? Given that a $(1 - 1/e)$ -approximation for welfare maximization in combinatorial auctions (without truthfulness) was also discovered first for coverage functions [8], then for weighted sums of matroid rank functions [4] and later extended to monotone submodular functions [25], it seems reasonable to conjecture that the same might happen for truthful-in-expectation mechanisms.

Our results: We prove that this is not the case, and there is a significant separation between the class of coverage functions and general monotone submodular functions. More precisely, there is no truthful-in-expectation mechanism (even $(1 - \epsilon)$ -approximately truthful-in-expectation) for submodular combinatorial auctions in the value oracle model, guaranteeing an approximation better than $1/m^\gamma$ for some fixed $\epsilon, \gamma > 0$ (Theorem 5.1). In particular, the results of [11] cannot be extended to all monotone submodular functions. We also prove a similar result for the *flexible submodular combinatorial public projects* problem (see Section 4 for a history of this problem): there is no $(1 - \epsilon)$ -

approximately truthful-in-expectation mechanism providing approximation better than $1/m^\gamma$ for some $\gamma > 0$. This is true even in the case of a single player. The combinatorial public projects problem admits a simpler structure than combinatorial auctions, and hence we deal with it first to demonstrate our approach.

Our techniques: Our hardness results are obtained by combining two recently developed techniques: the *symmetry gap* technique for submodular functions [26], and the *direct hardness* approach for combinatorial auctions [9].

First, we consider the possibility of maximal-in-distributional range (MIDR) mechanisms. We endeavor to explain why the approach of [11] breaks down when applied to monotone submodular functions. The answer lies in a certain convexity phenomenon that can be exploited in a symmetry gap argument. The symmetry gap argument on its own rules out the approach of [11]. Furthermore, it is possible to generalize the argument to an arbitrary MIDR mechanism, and moreover amplify the gap to some constant power of m . In fact our approach rules out even non-uniform approximately-MIDR mechanisms.

In the case of combinatorial public projects (CPP), we prove that if non-uniformity is allowed, then approximately truthful-in-expectation mechanisms are no more powerful than MIDR mechanisms. Therefore, by ruling out MIDR mechanisms, we also rule out truthful-in-expectation mechanisms. In the case of combinatorial auctions, no such equivalence in power between truthful-in-expectation and MIDR mechanisms is known. Instead, we apply the direct hardness approach of Dobzinski [9] to identify a single player for whom the allocation problem in some sense mimics the CPP problem. Again, the symmetry gap argument can be used here, though payments complicate the picture. We address this difficulty by employing a scaling argument and invoking the separating hyperplane theorem — this allows us to essentially get rid of the payments and use the same gap amplification technique we used for the CPP problem.

Organization: After the preliminaries (Section 2), we present our intuition on the separation between coverage and submodular functions in Section 3. Then we present an overview of our proofs for combinatorial public projects (Section 4) and for combinatorial auctions (Section 5). The complete proofs are deferred to the full version.

2. PRELIMINARIES

2.1. Mechanism Design Basics

Mechanism Design Problems: We consider mechanism design problems where there are n players, and a set Ω of feasible solutions. Each player i has a non-negative *valuation function* $v_i : \Omega \rightarrow \mathbb{R}_+$. We are concerned with *welfare maximization* problems, where the objective is $\sum_{i=1}^n v_i(\omega)$.

Mechanisms: We consider direct-revelation mechanisms for mechanism design problems. Such a mechanism comprises an *allocation rule* \mathcal{A} , which is a function from (hopefully truthfully) reported valuation functions $v = (v_1, \dots, v_n)$ to an outcome $\mathcal{A}(v) \in \Omega$, and a *payment rule* p , which is a function from reported valuation functions to a required payment $p_i(v)$ from each player i . We allow the allocation and payment rules to be randomized. We restrict our attention to mechanisms that are individually rational in expectation — i.e. $\mathbf{E}[v_i(\mathcal{A}(v)) - p_i(v)] \geq 0$ — and the payments are non-negative in expectation — i.e. $\mathbf{E}[p_i(v)] \geq 0$ — for each player i and each input $v = (v_1, \dots, v_n)$, when the expectations are over the random coins of the mechanism.

Truthfulness: A mechanism with allocation and payment rules \mathcal{A} and p is *truthful-in-expectation* if every player always maximizes its expected payoff by truthfully reporting its valuation function, meaning that

$$\mathbf{E}[v_i(\mathcal{A}(v)) - p_i(v)] \geq \mathbf{E}[v_i(\mathcal{A}(v'_i, v_{-i})) - p_i(v'_i, v_{-i})] \quad (1)$$

for every player i , (true) valuation function v_i , (reported) valuation function v'_i , and (reported) valuation functions v_{-i} of the other players. The expectation in (1) is over the coin flips of the mechanism. If (1) holds for every flip of the coins, rather than merely in expectation, we call the mechanism *universally truthful*.

VCG-Based Mechanisms: Mechanisms for welfare maximization problems are often variants of the classical VCG mechanism. Recall that the *VCG mechanism* is defined by the (generally intractable) allocation rule that selects the welfare-maximizing outcome with respect to the reported valuation functions, and the payment rule that charges each player i a bid-independent “pivot term” minus the reported welfare earned by other players in the selected outcome. This (deterministic) mechanism is truthful; see e.g. [22].

Let $\text{dist}(\Omega)$ denote the probability distributions over the set of feasible solutions Ω , and let $\mathcal{R} \subseteq \text{dist}(\Omega)$ be a compact subset of them. The corresponding *Maximal in Distributional Range (MIDR)* allocation rule is defined as follows: given reported valuation functions v_1, \dots, v_n , return an outcome that is sampled randomly from a distribution $D^* \in \mathcal{R}$ that maximizes the expected welfare $\mathbf{E}_{\omega \sim D}[\sum_i v_i(\omega)]$ over all distributions $D \in \mathcal{R}$. Analogous to the VCG mechanism, there is a (randomized) payment rule that can be coupled with this allocation rule to yield a truthful-in-expectation mechanism (see [6]). We note that

deterministic MIDR allocation rules — i.e. those where \mathcal{R} is a set of point distributions — are called *maximal-in-range (MIR)*.

Approximate Truthfulness: For $\epsilon \geq 0$, a mechanism with allocation and payment rules \mathcal{A} and p is $(1 - \epsilon)$ -*approximately truthful-in-expectation* if

$$\mathbf{E}[v_i(\mathcal{A}(v)) - p_i(v)] \geq (1 - \epsilon)\mathbf{E}[v_i(\mathcal{A}(v'_i, v_{-i})) - p_i(v'_i, v_{-i})] \quad (2)$$

for every player i , (true) valuation function v_i , (reported) valuation function v'_i , and (reported) valuation functions v_{-i} of the other players. The expectation in (2) is over the coin flips of the mechanism. Using the fact that payments are non-negative in expectation, a $(1 - \epsilon)$ -approximately truthful-in-expectation mechanism also satisfies the following weaker condition. (This condition is sufficient for our hardness results.)

$$\mathbf{E}[v_i(\mathcal{A}(v)) - p_i(v)] \geq \mathbf{E}[(1 - \epsilon)v_i(\mathcal{A}(v'_i, v_{-i})) - p_i(v'_i, v_{-i})] \quad (3)$$

Approximately truthful mechanisms are related to *approximately maximal-in-distributional-range* allocation rules. An allocation rule $\mathcal{A} : \mathcal{V} \rightarrow \Omega$ is $(1 - \epsilon)$ -approximately maximal-in-distributional range if it fixes a $\mathcal{R} \subseteq \text{dist}(\Omega)$, and returns an outcome that is sampled from $D^* \in \mathcal{R}$ that $(1 - \epsilon)$ -approximately maximizes the expected welfare $\mathbf{E}_{\omega \sim D}[\sum_i v_i(\omega)]$ over all distributions $D \in \mathcal{R}$. We show in Appendix A a sense in which approximately maximal-in-distributional-range allocation rules are no less powerful — in terms of approximating the social welfare — than approximately truthful-in-expectation mechanisms.

Our main reason for considering the notion of approximate truthfulness is that the mechanisms of [11], [13], if implemented in the value oracle model, are only approximately truthful-in-expectation (for an arbitrarily small $\epsilon > 0$, see [12]). The value oracle model seems too weak to make the mechanisms of [11], [13] exactly truthful-in-expectation; however, [12] makes it quite conceivable that there might be an approximately truthful-in-expectation mechanism for combinatorial auctions with submodular valuations.

2.2. Combinatorial Auctions

In *Combinatorial Auctions* there is a set M of m items, and a set of n players. Each player i has a valuation function $v_i : 2^M \rightarrow \mathbb{R}_+$ that is normalized ($v_i(\emptyset) = 0$) and monotone ($v_i(A) \leq v_i(B)$ whenever $A \subseteq B$). A feasible solution is an *allocation* (S_1, \dots, S_n) , where S_i denotes the items assigned to player i , and $\{S_i\}_i$ are mutually disjoint subsets of M . Player i 's value for outcome (S_1, \dots, S_n) is equal to $v_i(S_i)$. The goal is to choose an allocation maximizing *social welfare*: $\sum_i v_i(S_i)$.

2.3. Combinatorial Public Projects

In *Combinatorial Public Projects* there is a set $[m] = \{1, \dots, m\}$ of *projects*, a cardinality bound k such that

$0 \leq k \leq m$, and a set $[n] = \{1, \dots, n\}$ of *players*. Each player i has a valuation function $v_i : 2^{[m]} \rightarrow \mathbb{R}_+$ that is normalized ($v_i(\emptyset) = 0$) and monotone ($v_i(A) \leq v_i(B)$ whenever $A \subseteq B$). In this paper, we focus on the *flexible* variant of combinatorial public projects: a feasible solution is a set $S \subseteq [m]$ of projects with $|S| \leq k$. Player i 's value for outcome S is equal to $v_i(S)$. Prior work [24], [3], [9] has also considered the *exact* variant, where a feasible solution is a set $S \subseteq [m]$ with $|S| = k$. In both variants, the goal is to choose a feasible set S maximizing *social welfare*: $\sum_i v_i(S)$.

3. INTUITION

- WHAT FAILS FOR SUBMODULAR FUNCTIONS?

The main obstacle in proving our hardness result for submodular functions is the fact that the natural subclass of coverage functions *does* admit a truthful-in-expectation $(1 - 1/e)$ -approximation [11]. In the absence of strategic considerations, coverage functions capture the full difficulty of submodular functions in the context of welfare maximization, in the sense that they exhibit the same hardness threshold of $1 - 1/e$. Hence, it is not immediately clear where the dramatic jump in hardness should come from.

Let us recall the main idea of [11]: Let $f : 2^M \rightarrow \mathbb{R}_+$ be a submodular set function. Given $\mathbf{x} \in [0, 1]^M$, the expected value of $f(S)$ when S includes each item j independently with probability x_j is measured by the *multilinear extension* $F(\mathbf{x})$, which has been previously used in work on submodular maximization [4], [25], [17], [26], [23]. F is an *extension* of f , in the sense that it agrees with f on integer points, and therefore maximizing $F(\mathbf{x})$ over fractional allocations would yield an optimal algorithm. However, $F(\mathbf{x})$ is not a concave function and can be maximized only approximately. Instead, the authors of [11] consider a different rounding process — which they call the *Poisson rounding scheme* — that includes each j in S with probability $1 - e^{-x_j}$ instead. The expected value of applying the Poisson rounding scheme to a point \mathbf{x} is measured by a modified function $F^{exp}(x_1, \dots, x_m) = F(1 - e^{-x_1}, \dots, 1 - e^{-x_m})$, which fortuitously *turns out to be concave* for a subclass of submodular functions, including coverage functions and weighted sums of matroid rank functions. In this case, $F^{exp}(\mathbf{x})$ can be maximized exactly, and yields a maximal-in-distributional-range algorithm whose range is the image of the Poisson rounding scheme. Since the ratio between $F(\mathbf{x})$ and $F^{exp}(\mathbf{x})$ is bounded by $1 - 1/e$, this leads to a truthful-in-expectation $(1 - 1/e)$ -approximation.

The first question is whether F^{exp} can be maximized for any monotone submodular function. It was observed by the authors of [11] that F^{exp} is not concave for every submodular function: one example is the budget-additive function $f(S) = \min\{\sum_{i \in S} w_i, 2\}$ where $w_1 = w_2 = w_3 = 1$ and $w_4 = 2$. Hence convex optimization techniques cannot be used for $F^{exp}(\mathbf{x})$ directly; still, perhaps $F^{exp}(\mathbf{x})$ could

be maximized for a different reason. We prove that this is impossible, using a *symmetry gap* argument [15], [20], [26].

The budget-additive function above does not lend itself well to the symmetry gap argument, because there is a clear asymmetry between the elements of weight 1 and the element of weight 2. Instead, we construct an example where F^{exp} is not concave and all elements are in some sense “equivalent”. For this purpose, we use the following construction: If $f_1, f_2 : 2^M \rightarrow [0, 1]$ are monotone submodular functions, then

$$f(S) = 1 - (1 - f_1(S))(1 - f_2(S))$$

is also a monotone submodular function (we omit the proof). In particular, let $M = M_1 \cup M_2$, $|M_1| = |M_2| = m$, $|M| = 2m$, and let $f_i(S) = \min\{\frac{1}{\alpha m} |S \cap M_i|, 1\}$ for some $\alpha > 0$. These are budget-additive and hence monotone submodular functions. Therefore the following function is also monotone submodular:

$$\begin{aligned} f(S) &= 1 - (1 - f_1(S))(1 - f_2(S)) \\ &= 1 - \left(1 - \frac{1}{\alpha m} |S \cap M_1|\right)_+ \left(1 - \frac{1}{\alpha m} |S \cap M_2|\right)_+ \end{aligned}$$

Here, $(y)_+ = \max\{y, 0\}$ denotes the positive part of a number. Let's consider the function $F^{exp}(x_1, \dots, x_{2m}) = F(1 - e^{-x_1}, \dots, 1 - e^{-x_{2m}})$. If $m \rightarrow \infty$, a random set obtained by sampling with probabilities $1 - e^{-x_i}$ will have cardinality very close to $\sum(1 - e^{-x_i})$. We obtain

$$\begin{aligned} F^{exp}(\mathbf{x}) &\simeq 1 - \left(1 - \frac{1}{\alpha m} \sum_{i \in M_1} (1 - e^{-x_i})\right)_+ \\ &\quad \cdot \left(1 - \frac{1}{\alpha m} \sum_{j \in M_2} (1 - e^{-x_j})\right)_+ \end{aligned}$$

The reader can verify that this function is concave for $\alpha = 1$. But this is a very special coincidence. (The reason is that f for $\alpha = 1$ can be represented as a coverage function.) Any smaller value of α , for instance $\alpha = 1/2$, gives a non-concave function F^{exp} , as can be seen by checking $\mathbf{x} = \mathbf{1}_{M_1}$, $\mathbf{x} = \mathbf{1}_{M_2}$ and $\mathbf{x} = \frac{1}{2} \mathbf{1}_M$: $F^{exp}(\mathbf{1}_{M_1}) = F^{exp}(\mathbf{1}_{M_2}) = 1 - (-1 + 2e^{-1})_+ = 1$ (note that $-1 + 2e^{-1} < 0$), while the value at the midpoint is $F^{exp}(\frac{1}{2} \mathbf{1}_M) \simeq 1 - (-1 + 2e^{-1/2})^2 = 4e^{-1/2} - 4e^{-1} \simeq 0.955$. Therefore, we have an example where $F^{exp}(\mathbf{x})$ is not concave and moreover, all elements play the same symmetric role in f . (Formally, f has an element-transitive group of symmetries.) Functions of this type will play a crucial role in our proof.

The symmetry gap argument: The symmetry gap argument from [26], building up on previous work [15], [20], shows the following: Instances exhibiting some kind of symmetry can be blown up and modified in such a way that the only solutions that an algorithm can find (using a polynomial number of value queries) are symmetric with

respect to the same notion of symmetry. Thus the gap between symmetric and asymmetric solutions implies an inapproximability threshold. We use this argument here as follows. The instance above (for $\alpha = 1/2$) can be slightly modified as in [15], [20], [26], in such a way that it is impossible to find any solution that is asymmetric with respect to M_1, M_2 . Consider the optimization problem

$$\max\{F^{exp}(\mathbf{x}) : \sum x_i \leq m\}.$$

The best symmetric solution is $F^{exp}(\frac{1}{2}\mathbf{1}_M) \simeq 0.955$, while the optimum is $F^{exp}(\mathbf{1}_{M_1}) = 1$. The only solutions found by a polynomial number of value queries are the symmetric ones, and hence we cannot solve the optimization problem within a factor better than 0.955. A similar argument shows that we cannot solve the welfare maximization problem (for 2 players) with respect to $F^{exp}(\mathbf{x})$ within a factor better than 0.955.

In the following, we harness this construction towards showing that there can be no good maximum-in-distributional-range mechanism, and eventually, no good truthful-in-expectation mechanism.

4. HARDNESS FOR COMBINATORIAL PUBLIC PROJECTS

We start with the combinatorial public project problem. The (exact) combinatorial public project problem was introduced in [24] as a model problem for the study of truthful approximation mechanisms. This problem is better understood than combinatorial auctions, in the sense that a useful characterization of all deterministic truthful mechanisms is known: every truthful mechanism for 2 players is an *affine maximizer* — a weighted generalization of maximal-in-range mechanisms [24]. Using this characterization, it was proved in [24] that the exact submodular CPP problem does not admit any (deterministic) truthful $m^{\epsilon-1/2}$ -approximation using a subexponential amount of communication, and moreover there is no $m^{\epsilon-1/2}$ -approximation even for a certain class of succinctly represented submodular valuations unless $NP \subseteq BPP$. In contrast, the simple greedy algorithm is a non-truthful $(1 - 1/e)$ -approximation algorithm for this problem [21]. This was the first example of such a dramatic gap in approximability between truthful mechanisms and non-truthful algorithms.

In follow-up work, a simpler characterization-type statement for CPP was shown in [3]: Every truthful mechanism for a single player with a coverage valuation can, via a non-uniform polynomial time reduction, be converted to a truthful maximal-in-range mechanism without degrading its approximation ratio. Since every truthful mechanism for n players must embed a truthful mechanism for a single player, this allowed the authors to restrict attention to maximal-in-range mechanisms for a single player in proving an $m^{\epsilon-1/2}$ -approximation threshold for CPP with coverage valuations, assuming that $NP \not\subseteq P/poly$. The following easy converse of their characterization is notable: A maximal-in-range

mechanism for CPP with a single player can directly be used as a maximal-in-range mechanism for any number of players.

Recently, it was proved by Dobzinski [9] that the exact variant of the submodular CPP problem (under the constraint $|S| = k$) does not admit a truthful-in-expectation $m^{\epsilon-1/2}$ -approximation in the value oracle model. However, as noted in [9], the flexible variant of CPP (under the constraint $|S| \leq k$) is arguably more natural in the strategic setting. For the flexible variant of CPP, [9] proves that there is no universally truthful $m^{\epsilon-1/2}$ -approximation, but leaves open the possibility of a better truthful-in-expectation mechanism. Problems that have a packing structure like flexible CPP have historically proven to be easier to approximate using truthful-in-expectation mechanisms [19], [6], [10], [11]. Flexible CPP has exhibited a similar pattern; Dughmi [13] recently designed a truthful-in-expectation $(1 - 1/e)$ -approximation mechanism for CPP when players have explicit coverage valuations (which is optimal regardless of strategic issues [14]), and more generally when players have matroid rank sum valuations that support a certain randomized variant of value queries.

Transformation to MIDR mechanisms: While deterministic truthful mechanisms for the CPP problem are no more powerful in terms of approximation than maximal-in-range mechanisms [24], [3], the situation is slightly more complicated for randomized mechanisms. It is not clear whether truthful-in-expectation mechanisms are equivalent to maximal-in-distributional-range mechanisms. Nonetheless, we prove the following.

Theorem 4.1. *For every $\epsilon \geq 0$ and $c(m) > 0$ the following holds. If there is a $(1 - \epsilon)$ -approximately truthful-in-expectation mechanism \mathcal{M} for the (exact or flexible) CPP problem that achieves a $c(m)$ -approximation for submodular valuations on m elements, then for any $\delta > 0$ there is a non-uniform $(1 - 3\epsilon - \delta)$ -approximately maximal-in-distributional-range mechanism \mathcal{M}' that achieves a $c(m)$ -approximation for submodular valuations on m elements and uses at most m more value queries than \mathcal{M} .*

By a non-uniform mechanism, we mean a separate fixed mechanism for each input size m ; i.e., the size of the program can depend arbitrarily on m . The only bound on the non-uniform mechanism is the number of value queries used. The main idea is that although the range of prices offered by a truthful-in-expectation mechanism can be unbounded, the mechanism can be made MIDR “in the limit”, when the input valuation is scaled by a sufficiently large constant. This constant can be fixed for each input size m and acts as an “advice string” to the mechanism. We present the proof in the Appendix.

Hardness for MIDR mechanisms: Our hardness result for flexible submodular CPP rules out mechanisms purely based on the number of value queries used, and hence it

rules out even the non-uniform mechanisms mentioned in Theorem 4.1.

Theorem 4.2. *There are absolute constants $\epsilon, \gamma > 0$ such that there is no $(1 - \epsilon)$ -approximately maximal-in-distributional-range mechanism for the flexible submodular CPP problem with 1 player in the value oracle model, $\max\{f(S) : |S| \leq k\}$, achieving a better than $1/m^\gamma$ -approximation in expectation in the objective function, where m is the size of the ground set. This holds even for non-uniform mechanisms of arbitrary computational complexity, as long as the number of value queries is bounded by $\text{poly}(m)$.*

In the following, we present a sketch of the proof of this theorem. The full proof is deferred to the journal version.

Proof strategy: We assume that a mechanism optimizes over a range of distributions \mathcal{R} . (We assume for simplicity that the mechanism is MIDR rather than approximately MIDR.) We emphasize that the range \mathcal{R} is fixed beforehand, and the mechanism must optimize over \mathcal{R} for any particular submodular function f . This gives us a lot of flexibility in arguing about the properties of \mathcal{R} .

Suppose that the size of the ground set is $m = 2^{O(\ell)}$ and the cardinality bound is $k = m/2^\ell$. We consider $\ell + 1$ different “levels” of valuation functions. (See Figure 2.) At level 0, we have a set $A^{(0)}$ of $m/2^\ell$ items, where the valuation function is nonzero and additive. Assuming that the mechanism achieves a c -approximation, there must be a distribution $D_0 \in \mathcal{R}$ which allocates at least a c -fraction of $A^{(0)}$ in expectation to player i . This must be true for every set $A^{(0)}$ of size $m/2^\ell$. It will be useful to think of this set as random (and hidden from the mechanism.)

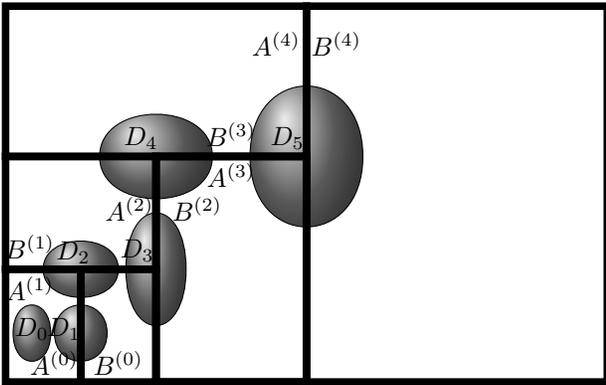


Figure 2. A bisection sequence $(A^{(j)}, B^{(j)})$, with the distributions D_j returned by the mechanism at level j . The density of D_j increases in a certain technical sense exponentially in j , although much slower than 2^j .

At level j , $1 \leq j \leq \ell$, we have a (random) set $A^{(j)}$ of $m/2^{\ell-j}$ items, which is partitioned randomly into two sets $A^{(j-1)} \cup B^{(j-1)}$ of equal size; these are level- $(j-1)$ sets. The valuation function at level j will be as in Section 3 but restricted to the set $A^{(j)} = A^{(j-1)} \cup B^{(j-1)}$ (the two parts

play the role of M_1, M_2 from Section 3). The mechanism can detect the set $A^{(j)}$; however, the partition of $A^{(j)}$ into $A^{(j-1)} \cup B^{(j-1)}$ remains hidden. By the symmetry gap argument, the mechanism cannot learn what the partition is, and hence any distribution D_j returned by the algorithm will be with high probability balanced with respect to $(A^{(j-1)}, B^{(j-1)})$. The MIDR property implies that this distribution must be “dense” enough in order to beat the distribution D_{j-1} guaranteed by the previous level, which is sensitive to the partition $(A^{(j-1)}, B^{(j-1)})$. (By density, we mean a certain notion of average size for sets sampled from D_j .) Since distributions concentrated inside $A^{(j-1)}$ or $B^{(j-1)}$ are more profitable than distributions balanced between $(A^{(j-1)}, B^{(j-1)})$, we will ideally obtain a constant-factor boost in density at each level. As ℓ grows, this will eventually contradict the fact that the mechanism cannot choose more than k items.

Finding the right definition of density that yields a constant-factor boost at each level is the main technical difficulty. The most natural definition of density seems to be the expected size of the set returned by the mechanism. However, this notion does not yield the desired boost. (This is related to the fact that we cannot get any contradiction for coverage functions.) The notion of density that turns out to be useful is more complicated; it is derived from functions that exhibit non-concave behavior of the extension F^{exp} . This strategy will be made more explicit in the following.

The symmetry gap: At level $j+1$, we consider valuation functions of the form

$$f_{A^{(j)}, B^{(j)}}(S) = 1 - \left(1 - \phi\left(\frac{|S \cap A^{(j)}|}{|A^{(j)}|}\right)\right) \cdot \left(1 - \phi\left(\frac{|S \cap B^{(j)}|}{|B^{(j)}|}\right)\right)$$

where $\phi : [0, 1] \rightarrow [0, 1]$ is a suitable non-decreasing concave function. Note that under this valuation function, the value of a (random) set R depends only on how many elements it takes from $A^{(j)}$ and $B^{(j)}$. In particular, if we denote $X_j = \frac{|R \cap A^{(j)}|}{|A^{(j)}|}$, $Y_j = \frac{|R \cap B^{(j)}|}{|B^{(j)}|}$, then we have

$$\mathbf{E}[f_{A^{(j)}, B^{(j)}}(R)] = \mathbf{E}[1 - (1 - \phi(X_j))(1 - \phi(Y_j))].$$

Since this expected value depends only on X_j, Y_j , we say that X_j, Y_j represent the distribution of R .

By the symmetry gap argument [15], [26] (if the valuation function is suitably perturbed and the partition $(A^{(j)}, B^{(j)})$ is random), the mechanism with high probability returns a solution $R^{(j+1)}$ independent of the partition and hence symmetric with respect to it. Denoting $X_{j+1} = \frac{|R^{(j+1)} \cap A^{(j+1)}|}{|A^{(j+1)}|}$, we obtain that the mechanism returns expected value

$$\mathbf{E}[f_{A^{(j)}, B^{(j)}}(R^{(j+1)})] = \mathbf{E}[1 - (1 - \phi(X_{j+1}))^2]$$

which is typically less than $\mathbf{E}[1 - (1 - \phi(X_j))(1 - \phi(Y_j))]$ if $X_{j+1} = \frac{1}{2}(X_j + Y_j)$ and $X_j \neq Y_j$. (There are certain

error terms arising from the symmetry gap argument but let us ignore them here.) The main point here is that if the mechanism is MIDR, then the expected value of the returned random set $\mathbf{E}[f_{A^{(j)}, B^{(j)}}(R^{(j+1)})]$ must be at least that of any other random set whose distribution is in the range - in particular, the random set $R^{(j)}$ whose presence in the range we prove at the previous level. If this random set $R^{(j)}$ is represented by the random variables X_j, Y_j , then the mechanism must return a distribution represented by X_{j+1} such that

$$\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] \geq \mathbf{E}[1 - (1 - \phi(X_j))(1 - \phi(Y_j))].$$

In fact we ignore the contribution of Y_j and use the weaker inequality

$$\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] \geq \mathbf{E}[\phi(X_j)]. \quad (4)$$

Hence, the existence of certain distributions in the range forces the existence of other distributions, satisfying the bound (4).

Gap amplification: Now we would like to say that if two distributions represented by X_j and X_{j+1} satisfy (4), then the distribution at level $j+1$ is “more dense” than the one at level j . Considering the scaling at different levels, we want to prove that X_{j+1} is “significantly larger” than $\frac{1}{2}X_j$. This is intuitive, since $X_{j+1} = \frac{1}{2}X_j$ is not enough to satisfy (4), for example when ϕ is linear. Unfortunately, (4) does not imply any useful relationship between the expectations $\mathbf{E}[X_j], \mathbf{E}[X_{j+1}]$, beyond $\mathbf{E}[X_{j+1}] \geq \frac{1}{2}\mathbf{E}[X_j]$. For example, we could have $X_j = 1$ with probability $\xi - \xi^2$ and 0 otherwise. Then $X_{j+1} = \frac{1}{2}\xi$ satisfies (4) for any concave function $\phi : [0, 1] \rightarrow [0, 1]$. This does not provide a constant-factor improvement over $\frac{1}{2}\mathbf{E}[X_j]$.

We still want to prove that X_{j+1} is in some sense “significantly larger” than $\frac{1}{2}X_j$. Our main technical inequality formalizing this intuition is the following: Define $\phi_\alpha(t) = \min\{\frac{t}{\alpha}, 1\}$. Then for any distribution in the range represented by X_j at level j and any $\alpha_j \in [0, 1]$, there is a distribution in the range represented by X_{j+1} at level $j+1$, and $\alpha_{j+1} \in [0, 1]$ such that

$$\alpha_{j+1}\mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})]^{1+\delta} \geq \frac{1+\delta^2}{2}\alpha_j\mathbf{E}[\phi_{\alpha_j}(X_j)]^{1+\delta} \quad (5)$$

where $\delta > 0$ is some (small) absolute constant. The use of $1 + \delta$ in the exponent is crucial here. The proof of (5) is quite technical; a rough sketch goes as follows.

By (4), there is a distribution at level $j+1$ such that $\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] \geq \mathbf{E}[\phi(X_j)]$. We split into 2 cases:

- 1) If $\Pr[X_{j+1} > \alpha_j\sqrt{\delta}]$ is at least some small constant, then we set $\alpha_{j+1} = \frac{1+\delta}{2}\alpha_j$. We use the fact that if $X_{j+1} > \alpha_j\sqrt{\delta}$, then $2X_{j+1}/\alpha_j$ is a constant-factor larger than $1 - (1 - X_{j+1}/\alpha_j)^2$, and therefore $\mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})]$ is a constant-factor larger than $\frac{1}{1+\delta}\mathbf{E}[\phi_{\alpha_j}(X_j)]$. From here, we prove $\alpha_{j+1}\mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})]^{1+\delta} \geq \frac{1+\delta^2}{2}\alpha_j\mathbf{E}[\phi_{\alpha_j}(X_j)]^{1+\delta}$.

- 2) If X_{j+1} is mostly in the interval $[0, \alpha_j\sqrt{\delta}]$, then we set $\alpha_{j+1} = \alpha_j\sqrt{\delta}$, i.e. much smaller than α_j . We gain in this case, because $\mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})]$ is much larger than $\mathbf{E}[\phi_{\alpha_j}(X_j)]$, by a factor of $\frac{1-O(\delta)}{2\sqrt{\delta}}$, and this factor is amplified by the power of $1 + \delta$. From this, we deduce $\alpha_{j+1}\mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})]^{1+\delta} \geq \frac{1+\delta^2}{2}\alpha_j\mathbf{E}[\phi_{\alpha_j}(X_j)]^{1+\delta}$.

The contradiction: Using (5), we arrive at a contradiction as follows. As we already mentioned, assuming that an MIDR mechanism provides a c -approximation for the CPP problem, then for any feasible set $A^{(0)}$ there must be a distribution D_0 in its range such that

$$\mathbf{E}[X_0] = \mathbf{E}_{R \sim D_0} \left[\frac{|R \cap A^{(0)}|}{|A^{(0)}|} \right] \geq c.$$

Now we apply the symmetry gap argument and the gap amplification technique to random pairs of sets $(A^{(j)}, B^{(j)})$ at each level j . Starting from $\mathbf{E}[X_0] \geq c$ and $\alpha_0 = 1$, by repeated use of (5) we obtain that there is $\alpha_\ell \in [0, 1]$ and a distribution at level ℓ represented by X_ℓ such that

$$\alpha_\ell(\mathbf{E}[\phi_{\alpha_\ell}(X_\ell)])^{1+\delta} \geq \left(\frac{1+\delta^2}{2}\right)^\ell c^{1+\delta}.$$

Note that $\alpha_\ell(\mathbf{E}[\phi_{\alpha_\ell}(X_\ell)])^{1+\delta} \leq \alpha_\ell\mathbf{E}[\phi_{\alpha_\ell}(X_\ell)] = \mathbf{E}[\min\{X_\ell, \alpha_\ell\}] \leq \mathbf{E}[X_\ell]$. So in fact

$$\mathbf{E}[X_\ell] \geq \left(\frac{1+\delta^2}{2}\right)^\ell c^{1+\delta} > \frac{2^{\delta^2\ell}}{2^\ell} c^{1+\delta}.$$

The meaning of X_ℓ is simply the fraction of the ground set that the mechanism returns at level ℓ . Since $m = 2^{O(\ell)}$, we have $2^{\delta^2\ell} \geq m^{(1+\delta)\gamma}$ for some constant $\gamma > 0$. If the approximation factor is $c \geq m^{-\gamma}$, then we get $\mathbf{E}[X_\ell] > 2^{-\ell}$, which would violate the cardinality constraint of the CPP problem.

5. HARDNESS FOR COMBINATORIAL AUCTIONS

The following is our main result for combinatorial auctions.

Theorem 5.1. *There are absolute constants $\epsilon, \gamma > 0$ such that there is no $(1-\epsilon)$ -approximately truthful-in-expectation mechanism for combinatorial auctions with monotone submodular valuation functions in the value oracle model, achieving a better than $1/n^\gamma$ -approximation in expectation in terms of social welfare, where the number of players is n and the number of items is $m = \text{poly}(n)$.*

Discussion: This theorem extends previous negative results for combinatorial auctions with submodular valuation functions, which were known only in the cases of deterministic truthful and randomized universally truthful mechanisms [9]. Also, it appears that as stated these results do not rule out approximately truthful mechanisms.

We remark that there is still the possibility of a truthful-in-expectation (TIE) mechanism in the “lottery-value” oracle

model which was introduced in [11]. Here, a player is able to provide the *exact* expectation $\mathbf{E}[v_i(\widehat{\mathbf{x}})]$ for a product distribution given by \mathbf{x} . Since the exact expectations $\mathbf{E}[v_i(\widehat{\mathbf{x}})]$ are hard to compute even in very special cases like the budget-additive case, this is a severe limitation. Our hardness result does not apply directly to this stronger oracle model. However, what our result implies is that if a truthful-in-expectation mechanism exists in the lottery-value model, then it must be very sensitive to the accuracy of the oracle’s answers, and does not remain even approximately truthful-in-expectation if the oracle’s answers involve some small noise. This is because if we had a mechanism in the lottery-value oracle model, which remains approximately TIE under small noise in the oracle and provides a good approximation, then we could simulate this mechanism in the value oracle model (by sample-average approximation). Thus we would obtain an approximately TIE mechanism contradicting Theorem 5.1.

Proof strategy: Our hardness result for combinatorial public projects (Section 4) can be adapted to show that there is no (approximately) MIDR mechanism for submodular combinatorial auctions that guarantees a good approximation ratio. However, unlike in CPP, we are unable to prove that truthful-in-expectation mechanisms and MIDR algorithms are equivalent in power (even in the approximate sense). This is not surprising, since randomized truthful mechanisms that are not maximal-in-distributional-range have been designed for combinatorial auctions (see for example [5]). Therefore, additional ideas are needed to rule out all truthful-in-expectation mechanisms. Such ideas have been recently put forth in a paper by Dobzinski [9]. The *direct hardness* approach of [9] provides a way to avoid the characterization step and instead attack the truthful mechanism directly. This idea applies to truthful-in-expectation mechanisms as well.

The main idea of the direct hardness approach can be stated as follows. If we identify a special player whose range of possible allocations is sufficiently “rich” when the valuations of other players are fixed to particular functions, then we can work with the special player directly using the *taxation principle*: There is a fixed price for each distribution over allocations in the “range” of the mechanism as the special player varies his valuation, and the mechanism outputs the distribution in this range that maximizes the player’s utility (his expected value for the distribution on allocations less the price of that distribution). Thus, our symmetry gap techniques from Section 4 apply here quite naturally, though the presence of payments poses an additional technical challenge that was not present for CPP. In the following, we present a sketch of our proof.

The basic instance: We start from the following “basic instance”. For an integer ℓ , we construct instances with $|N| = n = 2^\ell$ players and $|M| = m = \text{poly}(n)$ items. Each player has a “polar valuation” v_i^* (as in [9]), where items in a certain set $A_i^{(0)}$ have value 1 for player i and

other items have (small) value $\omega > 0$. The sets $A_i^{(0)}$ are chosen independently at random, under the constraint that $|A_i^{(0)}| = m/n$.

A counting argument shows that if a mechanism provides a c -approximation in social welfare, then there must be a player whose allocated set $R_i^{(0)}$ overlaps significantly with his desired set $A_i^{(0)}$:

$$\mathbf{E}[|R_i^{(0)} \cap A_i^{(0)}|] \geq (c/4 - \omega)\mathbf{E}[|R_i^{(0)} \cup A_i^{(0)}|].$$

By an averaging argument, this is also true for a certain fixed choice of the other players’ valuations. In the following, we fix that choice and consider varying valuations for player i only, who we refer to as the “special player”. We also drop the index i , since we do not consider the other players anymore.

In the following, we set $\omega = c/8$, so that $\mathbf{E}[|R^{(0)} \cap A^{(0)}|] \geq \omega\mathbf{E}[|R^{(0)} \cup A^{(0)}|]$. Hence we can estimate the expected value received by the special player as follows:

$$\begin{aligned} \mathbf{E}[v^*(R^{(0)})] &= \mathbf{E}[|R^{(0)} \cap A^{(0)}|] + \omega\mathbf{E}[|R^{(0)} \setminus A^{(0)}|] \\ &\leq 2\mathbf{E}[|R^{(0)} \cap A^{(0)}|]. \end{aligned}$$

Denoting $X_0 = \frac{|R^{(0)} \cap A^{(0)}|}{|A^{(0)}|}$, we have $\mathbf{E}[v^*(R^{(0)})] \leq \frac{2m}{n}\mathbf{E}[X_0]$. Also, $\mathbf{E}[v^*(R^{(0)})] \geq \mathbf{E}[|R^{(0)} \cap A^{(0)}|] = \frac{m}{n}\mathbf{E}[X_0]$. So the special player’s utility in the basic instance (in expectation over the random instances) is $\frac{m}{n}\mathbf{E}[X_0]$, up to a factor of 2.

Symmetry gap again: We consider valuations for the special player at ℓ levels, in the same form that we considered in the case of combinatorial public projects. The difference now is that the mechanism is not necessarily maximal-in-distributional-range. Instead, we use the definition of truthfulness in expectation directly. The same symmetry gap argument as in Section 4 gives the following: If there is a random set $R^{(j)}$ possibly allocated at level j at a price P_j , and $X_j = \frac{|R^{(j)} \cap A^{(j)}|}{|A^{(j)}|}$, then there is a random set $R^{(j+1)}$ possibly allocated at level $j+1$ at a price P_{j+1} , and $X_{j+1} = \frac{|R^{(j+1)} \cap A^{(j+1)}|}{|A^{(j+1)}|}$, so that

$$\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] - \mathbf{E}[P_{j+1}] \geq \mathbf{E}[\phi(X_j)] - \mathbf{E}[P_j].$$

Again, we are ignoring certain error terms and we are also ignoring the issue of approximate truthfulness. Using the fact that the valuation functions can be scaled arbitrarily and the mechanism must still be truthful in expectation, we obtain that for any $\lambda', \lambda'' \geq 0$, there is distribution possibly allocated at level $j+1$ such that

$$\lambda'\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] - \lambda''\mathbf{E}[P_{j+1}] \geq \lambda'\mathbf{E}[\phi(X_j)] - \lambda''\mathbf{E}[P_j]. \quad (6)$$

Convex hulls and the separation argument: Our goal is to eliminate the prices from the picture, so that we can use arguments similar to Section 4. For that purpose, it is convenient to pass to convex hulls as follows. We

define the *distribution menu* \mathcal{M}_j at level j to consist of all distributions of pairs of random variables (X_j, P_j) , such that $X_j = \frac{|R^{(j)} \cap A^{(j)}|}{|A^{(j)}|}$ for some random set $R^{(j)}$ allocated for a level- j valuation at a price P_j . Then we define the *closure* of a distribution menu, $\overline{\mathcal{M}}_j$, to be the topological closure of the convex hull of \mathcal{M}_j (in the sense of taking convex combinations of distributions). By convexity, (6) still holds in the sense that for any (X_j, P_j) with a distribution in $\overline{\mathcal{M}}_j$ and any $\lambda', \lambda'' \geq 0$, there is (X_{j+1}, P_{j+1}) with a distribution in $\overline{\mathcal{M}}_{j+1}$ such that (6) holds.

A convex separation argument, essentially Farkas' lemma in 2 dimensions, actually implies the following. For any (X_j, P_j) with a distribution in $\overline{\mathcal{M}}_j$, there is (X_{j+1}, P_{j+1}) with a distribution in $\overline{\mathcal{M}}_{j+1}$ such that $\mathbf{E}[P_{j+1}] \leq \mathbf{E}[P_j]$ and

$$\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] \geq \mathbf{E}[\phi(X_j)].$$

In other words, there is a distribution in the closure of the menu at level $j + 1$ at a price no higher than the price we had at level j , and the respective random variables X_j, X_{j+1} satisfy the same relationship (4) that we had in Section 4. The rest of the proof goes exactly as in Section 4, using (5) and eventually producing a distribution represented by (X_ℓ, P_ℓ) in $\overline{\mathcal{M}}_\ell$ such that $\mathbf{E}[P_\ell] \leq \mathbf{E}[P_0]$ and

$$\mathbf{E}[X_\ell] \geq \left(\frac{1 + \delta^2}{2}\right)^\ell (\mathbf{E}[X_0])^{1+\delta}.$$

Here, (X_0, P_0) represents the distribution and price allocated in the basic instance. Now we consider the utility that the distribution represented by (X_ℓ, P_ℓ) would provide in the basic instance: since every element has value at least ω there, the utility would be

$$\begin{aligned} \mathbf{E}[v^*(R^{(\ell)}) - P_\ell] &\geq \mathbf{E}[\omega m X_\ell - P_\ell] \\ &\geq \omega m \left(\frac{1 + \delta^2}{2}\right)^\ell (\mathbf{E}[X_0])^{1+\delta} - \mathbf{E}[P_0]. \end{aligned} \quad (7)$$

Recall that the distribution of (X_ℓ, P_ℓ) is not on the menu \mathcal{M}_ℓ but rather in its convex hull. However, by using the properties of the convex hull, there must be a distribution on the actual menu \mathcal{M}_ℓ that satisfies the same linear inequality. So we can assume without loss of generality that the distribution of (X_ℓ, P_ℓ) is on the actual menu at level ℓ , and $R^{(\ell)}$ is the respective random set that would be allocated to the special player if he declared a level- ℓ valuation.

Recall that in the basic instance, the value received by the special player is at most $\frac{2m}{n} \mathbf{E}[X_0]$, and the respective utility is at most $\frac{2m}{n} \mathbf{E}[X_0] - \mathbf{E}[P_0]$. We also have $n = 2^\ell$ and $\mathbf{E}[X_0] \geq c/4 - \omega = c/8$. If $c = 8\omega \geq n^{-\gamma}$ for a suitable constant $\gamma > 0$, we would obtain from (7) that the special player could substantially improve his utility in the basic instance by declaring a level- ℓ valuation instead. We conclude that this would contradict the property of truthfulness in expectation.

REFERENCES

- [1] N. Alon and J.H. Spencer. The probabilistic method (2nd edition). Wiley Interscience, 2000.
- [2] D. Buchfuhrer, S. Dughmi, H. Fu, R. Kleinberg, E. Mossel, C. Papadimitriou, M. Schapira, Y. Singer, and C. Umans. Inapproximability for VCG-based combinatorial auctions. *Proc. of 21th ACM-SIAM SODA*, 518–536, 2010.
- [3] D. Buchfuhrer, M. Schapira and Y. Singer. Computation and incentives in combinatorial public projects. *Proc. of 11th ACM EC*, 33–42, 2010.
- [4] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. *Proc. of 12th IPCO*, 182–196, 2007.
- [5] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. *Proc. of APPROX 2007*, 89–103.
- [6] S. Dobzinski and S. Dughmi. On the power of randomization in algorithmic mechanism design. *Proc. of 50th IEEE FOCS*, 505–514, 2009.
- [7] S. Dobzinski and N. Nisan. Limitations of VCG-based mechanisms. *Proc. of 39th ACM STOC*, 338–344, 2007.
- [8] S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. *Proc. of 17th SODA*, 1064–1073, 2006.
- [9] S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. *Proc. of 43rd ACM STOC*, 139–148, 2011.
- [10] S. Dughmi and T. Roughgarden. Black box randomized reductions in algorithmic mechanism design. *Proc. of 51st IEEE FOCS*, 775–784, 2010.
- [11] S. Dughmi, T. Roughgarden and Q. Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. *Proc. of 43rd ACM STOC*, 149–158, 2011.
- [12] S. Dughmi, T. Roughgarden, J. Vondrák and Q. Yan. An approximately truthful-in-expectation mechanism for combinatorial auctions using value queries. Manuscript, 2011.
- [13] S. Dughmi. A truthful randomized mechanism for combinatorial public projects via convex optimization. *Proc. of 12th ACM EC*, 263–272, 2011.
- [14] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [15] U. Feige, V. Mirrokni and J. Vondrák. Maximizing a non-monotone submodular function. *Proc. of 48th IEEE FOCS*, 461–471, 2007.
- [16] S. Khot, R. Lipton, E. Markakis and A. Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica* 52:1, 3–18, 2008.
- [17] A. Kulik, H. Shachnai and T. Tamir. Maximizing submodular set functions subject to multiple linear constraints. *Proc. of 20th ACM-SIAM SODA*, 545–554, 2009.
- [18] R. Lavi, A. Mu'alem and N. Nisan. Towards a characterization of truthful combinatorial auctions. *Proc. of 44th IEEE FOCS*, 574–583, 2003.
- [19] R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. *Proc. of 46th IEEE FOCS*, 595–604, 2005.
- [20] V. Mirrokni, M. Schapira and J. Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. *Proc. of 9th ACM EC*, 70–77, 2008.
- [21] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis

of approximations for maximizing submodular set functions - I. *Math. Prog.*, 14:265–294, 1978.

- [22] N. Nisan. Introduction to mechanism design (for computer scientists). *Algorithmic Game Theory* (N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, eds.), Cambridge University Press, 2007.
- [23] S. Oveis Gharan and J. Vondrák. Submodular maximization by simulated annealing. *Proc. of 22nd ACM-SIAM SODA*, 1098-1117, 2011.
- [24] C. Papadimitriou, M. Schapira and Y. Singer. On the hardness of being truthful. *Proc. of 49th IEEE FOCS*, 250–259, 2008.
- [25] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. *Proc. of 40th ACM STOC*, 67–74, 2008.
- [26] J. Vondrák. Symmetry and approximability of submodular maximization problems. *Proc. of 50th IEEE FOCS*, 251–270, 2009.

APPENDIX

In this appendix, we sketch the proof of Theorem 4.1. Let $\Omega = \{S \subseteq [m] : |S| \leq k\}$ be the set of outcomes of CPP. Let $\Delta(\Omega)$ denote the simplex in \mathbb{R}^Ω , representing the set of distributions over Ω . Let \mathcal{V} denote the set of submodular valuations on $[m]$. We think of \mathcal{V} as a subset of \mathbb{R}_+^Ω — specifically, each $v \in \mathcal{V}$ is a vector in \mathbb{R}_+^Ω , where v_S is the value of outcome S for a player with valuation v . We note that for each $v \in \mathcal{V}$, the infinity norm $\|v\|_\infty$ is equal to the value of the optimum solution. Fix ϵ, c, \mathcal{M} , and δ as in the statement of the theorem. Let $\mathcal{A} : \mathcal{V} \rightarrow \Delta(\Omega)$ be the allocation rule of \mathcal{M} when there is a single player. By assumption, \mathcal{A} is a c -approximation for $c > 0$ — specifically, $\frac{v^T \mathcal{A}(v)}{\|v\|_\infty} \geq c$.

We prove Theorem 4.1 by showing that there is a black-box reduction that converts \mathcal{A} to a new allocation rule \mathcal{B} that is $(1 - 3\epsilon - \delta)$ -MIDR. The reduction will be non-uniform — specifically, \mathcal{B} will utilize an advice string that depends on m , but is independent of the input valuation $v \in \mathcal{V}$. The length of the advice string will not be bounded, polynomially or otherwise — this is OK, since we are only interested in preserving value oracle lower-bounds. \mathcal{B} preserves the approximation ratio of \mathcal{A} , and moreover makes only m more value queries than does \mathcal{A} .

The proof consists of two main steps. First, we show that \mathcal{A} tends to a $(1 - \epsilon)$ -approximately maximal-in-distributional-range allocation rule “in the limit” as we scale up the valuations. Then, we use this fact to construct, via a non-uniform black box reduction, an allocation rule \mathcal{B} that approximates the limit behavior of \mathcal{A} , in the sense that it $(1 - 3\epsilon - \delta)$ -approximately maximizes over the range of \mathcal{A} . We briefly sketch the ideas in each of these steps next, deferring details to the full version of the paper.

Limit Behavior of the Mechanism: First, we make a trivial observation: if \mathcal{M} does not employ payments, then our assumption that it is $(1 - \epsilon)$ -approximately truthful-in-expectation implies that its allocation rule \mathcal{A} is $(1 - \epsilon)$ -MIDR. In the presence of payments, however, this no longer

follows. Our proof essentially shows that these payments grow slowly as valuations are scaled up. Specifically, if the mechanism \mathcal{M} charges a player with valuation $v \in \mathcal{V}$ a price $p(v)$, then if a scaling factor $\alpha > 1$ is applied to the player’s valuation, the ratio of the payment $p(\alpha v)$ to the player’s value $\alpha v^T \mathcal{A}(\alpha v)$ tends to zero as α tends to ∞ . Therefore, payments become “insignificant” for large valuations, and consequently \mathcal{A} tends to a $(1 - \epsilon)$ -MIDR allocation rule as valuations are scaled up. Our proof of this fact employs an analogue of weak monotonicity for approximate truthfulness, and we omit the details.

Approximating the Limit Behavior: Ideally, we would transform \mathcal{A} to an allocation rule that behaves as \mathcal{A} does in the limit. By the preceding discussion, such a “limit allocation rule” of \mathcal{A} would be $(1 - \epsilon)$ -MIDR. However, since our reduction must take finite time, we must settle for approximating the limit behavior of \mathcal{A} . Unfortunately, even that is non-trivial: given v , the ratio α by which we would need to scale v before coming close to the “limit” of $\mathcal{A}(\alpha v)$ is a complete mystery, and may be arbitrarily large. Therefore, we need to utilize some non-uniform advice to deduce that order of magnitude of the necessary scaling factor. An additional difficulty is that this advice must be independent of v — specifically, the advice may depend only on the number of items m .

We show that, for each fixed number of items m and $\delta > 0$, there is a $\tau = \tau(\delta)$ such that, for each $u \in \mathcal{V}$ with $\|u\|_\infty = 1$, the output $\mathcal{A}(\tau u)$ is within $1 - 2\epsilon - \delta$ of the limit behavior of \mathcal{A} — formally $u^T \mathcal{A}(\tau u) \geq (1 - 2\epsilon - \delta) \lim_{\alpha \rightarrow \infty} u^T \mathcal{A}(\alpha u)$. The number τ will serve as our advice string. We emphasize that τ is a *uniform* bound over all normalized “directions” $u \in \mathcal{V}$, and this uniformity is what allows the advice string to be independent of the valuation. The existence of such a uniform bound is a-priori not obvious, and the proof of this fact relies crucially on the approximate analogue of weak monotonicity, and employs the construction of a finite net of \mathcal{V} ; we omit the details.

Given $\tau = \tau(\delta)$ as a non-uniform advice string, it follows from the definition of τ and the realization that \mathcal{A} is $(1 - \epsilon)$ -MIDR in the limit that the following is an $(1 - 3\epsilon - \delta)$ -MIDR allocation rule: On input v , let $s(v)$ be a non-zero lower-bound on $\|v\|_\infty$, and output $\mathcal{A}(\tau \frac{v}{s(v)})$. Such a bound $s(v)$ is easily computed for submodular valuations using m value queries, by setting $s(v) = \max_j v(\{j\})$. Therefore, this completes our proof.