

Amplified Hardness of Approximation for VCG-Based Mechanisms

Shaddin Dughmi
Department of Computer Science
Stanford University
shaddin@cs.stanford.edu

Hu Fu
Department of Computer Science
Cornell University
hufu@cs.cornell.edu

Robert Kleinberg
Department of Computer Science
Cornell University
rdk@cs.cornell.edu

July 24, 2009

Abstract

If a two-player social welfare maximization problem does not admit a PTAS, we prove that any maximal-in-range truthful mechanism that runs in polynomial time cannot achieve an approximation factor better than $1/2$. Moreover, for the k -player version of the same problem, the hardness of approximation improves to $1/k$ under the same two-player hardness assumption. (We note that $1/k$ is achievable by a trivial deterministic maximal-in-range mechanism.) This hardness result encompasses not only deterministic maximal-in-range mechanisms, but also all universally-truthful randomized maximal in range algorithms, as well as a class of strictly more powerful truthful-in-expectation randomized mechanisms recently introduced by Dobzinski and Dughmi. Our result applies to any class of valuation functions that satisfies some minimal closure properties. These properties are satisfied by the valuation functions in all well-studied APX-hard social welfare maximization problems, such as coverage, submodular, and subadditive valuations.

We also prove a stronger result for universally-truthful maximal-in-range mechanisms. Namely, even for the class of budgeted additive valuations, which admits an FPTAS, no such mechanism can achieve an approximation factor better than $1/k$ in polynomial time.

1 Introduction

Do computational problems become harder when the inputs are supplied by selfish agents and the algorithm is required to operate in a way that incentivizes truth-telling? This question has been central to algorithmic mechanism design since the field’s inception [14]. The most famous positive result in the area is also one of the simplest: any efficient social-welfare-maximization algorithm can be transformed into a computationally efficient truthful mechanism using the celebrated VCG payment scheme [6, 11, 17]. It is also well known that this result does not extend to approximation algorithms: in order for an algorithm to be truthfully implemented by the VCG payment scheme, the algorithm must satisfy a property known as *maximal-in-range* (MIR) [15], which is unfortunately violated by most approximation algorithms. However, the technique of combining a maximal-in-range algorithm with the VCG payment scheme remains the only known general-purpose technique for designing truthful mechanisms for multi-parameter domains¹, and consequently a great deal of research has been devoted to searching for computationally efficient approximation algorithms that are maximal-in-range, e.g. [9, 10, 12], or proving hardness-of-approximation theorems for this class of algorithms, e.g. [8, 13, 16].

Combinatorial auctions are the most well-studied, and arguably the most important, class of mechanism design problems, and they furnish striking insights into the capabilities and limitations of maximal-in-range mechanisms. One can bound the approximation ratio of a combinatorial auction mechanism in terms of many parameters, including the number of players (henceforth denoted by k) or the number of items (henceforth, m). A particularly bleak picture emerges when one bounds the approximation ratio in terms of the number of players. Any combinatorial auction with k players has a trivial mechanism that simply packages all the items as a single bundle and awards this bundle to the bidder who values it most highly. This mechanism is computationally efficient, maximal in range, but only achieves approximation ratio $1/k$. Despite years of research on truthful combinatorial auctions, to our knowledge this dependence on k can be improved in only one combinatorial auction domain: the domain of multi-unit auctions. Here, the underlying social welfare maximization problem admits a deterministic FPTAS that is not maximal-in-range. Dobzinski and Nisan [9] discovered a deterministic 2-approximation that is MIR, and Dobzinski and Dughmi [7] discovered an FPTAS that outputs a randomized allocation satisfying a property called *maximal in distributional range* (MIDR). In fact, the Dobzinski-Dughmi mechanism satisfies a stronger property that we call *maximal in weighted range* (MIWR), meaning that the only use of randomization is to cancel the allocation (i.e., allocate no items) with some probability.

1.1 Our contributions

In this paper we show that there is an inherent reason why truthful mechanisms cannot break the “ $1/k$ barrier” for combinatorial auctions: *any approximation hardness at all in the underlying social welfare maximization problem is amplified to $(1/k + \epsilon)$ -hardness when one restricts the algorithm to be maximal in range*. In fact, our result extends to two classes of randomized MIDR mechanisms: those that choose an allocation deterministically and then toss coins to decide whether to cancel the allocation (MIWR mechanisms) and those that toss coins to choose a deterministic MIR mechanism and then run it (randomized MIR mechanisms), as well as the combination of the two. Our result applies to any class of valuations that is *regular*, meaning that it satisfies some natural closure

¹A multi-parameter domain is one in which an agent’s private information consists of more than just a single real-valued parameter.

properties to be specified in Section 2. These properties are satisfied by the valuation functions in all well-studied APX-hard social welfare maximization problems, such as coverage, submodular, XOS, and subadditive valuations.

Theorem 1.1. *Let \mathcal{C} be any regular class of valuations such that two-player social welfare maximization with valuations in \mathcal{C} does not admit a PTAS. Then for all $\epsilon > 0$, there is no polynomial-time randomized MIWR mechanism that achieves an approximation ratio greater than $1/k + \epsilon$ unless $\text{NP} \subseteq \text{P/poly}$.*

Even when social welfare maximization over \mathcal{C} admits an FPTAS, it may still be possible to prove that maximal-in-range mechanisms cannot improve on the trivial $1/k$ approximation factor. In fact, we are able to show this for an important class of valuations that admits an FPTAS: the class of *budgeted additive valuations*, in which each player i has a budget B_i , and her value for a bundle is equal to the sum of her values for the individual items, or to B_i , whichever is smaller. However, unlike Theorem 3.1, this hardness result is limited to universally-truthful randomized maximal-in-range mechanisms.

Theorem 1.2. *For all $\epsilon > 0$, no polynomial-time randomized MIR mechanism for combinatorial auctions with budgeted additive valuations can achieve an approximation ratio greater than $1/k + \epsilon$, unless $\text{NP} \subseteq \text{P/poly}$.*

1.2 Our techniques

A paradigm for proving hardness results of this sort was introduced by Papadimitriou, Schapira, and Singer in [16]. To prove that a certain approximation ratio cannot be achieved by maximal-in-range mechanisms, one proves that the underlying social welfare maximization problem exhibits a particularly strong form of self-reducibility: any maximal-in-range algorithm for optimizing over a sufficiently large subset of allocations can be transformed into an algorithm for optimizing over *all allocations* of a smaller set of items. The technical core of any such proof is a lemma showing that any sufficiently large range of allocations must “shatter” a fairly large subset S of the items, meaning that there is a set of players P such that all allocations of S to P occur in the range. In [16] the relevant shattering lemma was the famous Sauer-Shelah Lemma. But since the Sauer-Shelah Lemma is a statement about collections of a *subsets* of a ground set U , and the range of a combinatorial auction is a collection of *partial functions* from U to the set of players, we require new shattering lemmas that apply to partial functions.

Extending the Sauer-Shelah Lemma to partial functions is far from trivial: a lower bound on the cardinality of the range does not suffice to prove that it shatters a large set of items; for example, the set of all allocations that give a subset of the items to player 1 and no items to any other player constitutes an exponentially large range but does not shatter any nonempty set of items. Thus, one needs to carefully define what is meant by the hypothesis that the range is “large”, and also (in the case of more than two players) what is meant by the conclusion that it “shatters” a large set of items. In this paper, we provide two such lemmas. In both of them, U and V are finite sets with $|U| = m$, $|V| = k$, and R is a set of functions from U to $V \cup \{*\}$.

Lemma 1.3. *Suppose that for a random $f : U \rightarrow V$, with probability at least γ there is a $g \in R$ such that $g(x)$ differs from $f(x)$ on at most $\left(1 - \frac{q-1}{k} - \epsilon\right)m$ elements $x \in U$. Then there is a subset $S \subseteq U$ of cardinality at least δm (where $\delta > 0$ may depend on γ, ϵ, q, k) and a subset $T \subseteq V$ of cardinality q , such that every function from S to T occurs as the restriction of some $g \in R$.*

Lemma 1.4. *Suppose that ϵ, α, ℓ are constants such that $|R| > e^{\alpha m}$, and suppose that for every ℓ -tuple of functions $g_1, \dots, g_\ell \in R$, for some $1 \leq i < j \leq \ell$ there are at least ϵm elements $x \in U$ such that $g_i(x)$ and $g_j(x)$ are distinct elements of V . Then there is a subset $S \subseteq U$ of cardinality at least δm (where $\delta > 0$ may depend on $\epsilon, \alpha, \ell, k$) and a pair of elements $a, b \in V$, such that every function from S to $\{a, b\}$ occurs as the restriction of some $g \in R$.*

The first lemma, which underlies our proof of Theorem 3.1 and may be of independent interest, substitutes an assumption that R has small covering radius in the Hamming metric in place of the usual assumption that R has large cardinality. The second lemma, which underlies our proof of Theorem 1.2, generalizes and closely parallels a related lemma from [13]. We prove both lemmas in Appendix A.

To derive the lower bound for MIWR mechanisms, an additional idea is needed: rather than reducing directly converting an α -approximate MIR algorithm into an exact optimization algorithm over a smaller set of items, we convert an $(\alpha - \delta)$ -approximate MIWR algorithm into an $(\alpha + \delta)$ -approximate MIWR algorithm over a smaller set of items, and then we reach a contradiction by taking α to be the supremum of the approximation ratios achievable by polynomial-time MIWR mechanisms. Translating this idea into a rigorous proof requires a delicate induction over the number of players. Finally, to extend the result to randomized MIWR mechanisms, we show that any randomized MIWR mechanism can be transformed into a MIWR mechanism with polynomial advice, incurring a negligible loss in the approximation factor. The proof of this step closely parallels Adleman’s proof that $\text{BPP} \subseteq \text{P}/\text{poly}$.

1.3 History of these results

Our work builds on the work of [13], which obtained a weaker version of Theorem 1.2, also using the “shattering” technique. Independently and concurrently with our discovery of Theorem 1.2, a different proof of a similar result (limited to deterministic mechanisms, but obtaining optimal dependence on the number of items as well as players) was discovered by Buchfuhrer and Umans [5]. Our Lemma A.5, which constitutes a step in the proof of Theorem 3.1 and was discovered after we had read the proof of the Buchfuhrer-Umans result, uses a counting argument similar to their proof of a seemingly unrelated shattering lemma in [5].

2 Preliminaries

We assume the reader is familiar with standard terminology and notation regarding truthful mechanisms and approximation algorithms. Appendix D contains the relevant definitions.

2.1 Combinatorial Auctions

In a combinatorial auction there is a set $[m] = \{1, 2, \dots, m\}$ of items, and a set $[k] = \{1, 2, \dots, k\}$ 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 $A \subseteq B$).

An allocation of items M to the players N is a function $S : M \rightarrow N \cup \{*\}$. Notice that we do not require all items to be allocated. If an allocation S allocates all items – i.e. S maps M into N – we say S is a *total allocation*. The allocation that allocates no items is called the *empty allocation*.

For convenience, we use $S(j)$ to denote the player receiving item j , and we use S_i to denote the items allocated to player i . We use $\mathcal{X}(M, N)$ to denote the set of all allocations of M to N .

In combinatorial auctions, the feasible solutions are the allocations $\mathcal{X}([m], [k])$ of the items to the players. The *social welfare* of such an allocation S is defined as $\sum_i v_i(S_i)$. When the players have values $\{v_i\}_i$, we often use $v(S)$ as shorthand for the welfare of S . The goal in combinatorial auctions is to find an allocation that maximizes the social welfare.

2.2 Valuation Classes

The hardness of designing truthful combinatorial auction mechanisms depends on the allowable player valuations. Recall that a *valuation* over M is a function $v : 2^M \rightarrow \mathbb{R}^+$. We let \mathcal{V} denote the set of all valuations over all abstract finite sets M . A *valuation class* \mathcal{C} is a subset of \mathcal{V} . Examples of valuation classes include submodular valuations, subadditive valuations, single-minded valuations, etc. Our first result applies to any valuation class that satisfies some natural properties.

Definition 2.1. *We say a valuation class \mathcal{C} is regular if the following hold*

1. *Every valuation in \mathcal{C} is monotone and normalized.*
2. *The canonical valuation on any singleton set is in \mathcal{C} . Namely, for any item a the valuation $v : 2^{\{a\}} \rightarrow \mathbb{R}^+$, defined as $v(\{a\}) = 1$ and $v(\emptyset) = 0$, is in \mathcal{C} .*
3. *Closed under scaling: Let $v : 2^M \rightarrow \mathbb{R}^+$ be in \mathcal{C} , and let $c \geq 0$. The valuation $v' : 2^M \rightarrow \mathbb{R}^+$, defined as $v'(A) = c \cdot v(A)$ for all $A \subseteq M$, is also in \mathcal{C} .*
4. *Closed under disjoint union: Let M_1 and M_2 be disjoint sets. Let the valuations $v_1 : 2^{M_1} \rightarrow \mathbb{R}^+$ and $v_2 : 2^{M_2} \rightarrow \mathbb{R}^+$ be in \mathcal{C} . Their disjoint union $v_3 = v_1 \oplus v_2 : 2^{M_1 \cup M_2} \rightarrow \mathbb{R}^+$, defined as $v_3(A) = v_1(A \cap M_1) + v_2(A \cap M_2)$ for all $A \subseteq M_1 \cup M_2$, is in \mathcal{C} .*
5. *Closed under relabeling: Let M_1, M_2 be sets with a bijection $f : M_2 \rightarrow M_1$. If $v_1 : 2^{M_1} \rightarrow \mathbb{R}^+$ is in \mathcal{C} , then the valuation $v_2 : 2^{M_2} \rightarrow \mathbb{R}^+$ defined by $v_2(S) = v_1(f(S))$ is also in \mathcal{C} .*

Note that all regular valuation classes support *zero-extension*. More formally, let $M \subseteq M'$, and let $v : 2^M \rightarrow \mathbb{R}^+$ be in \mathcal{C} . The extension of v to M' , defined as $v'(A) = v(A \cap M)$ for all $A \subseteq M'$, is also in \mathcal{C} . In the context of combinatorial auctions, we use \mathcal{C}_m to denote the subset of valuation class \mathcal{C} that applies to items $[m]$.

Most well-studied valuation classes for which the underlying optimization problem is APX-hard are regular. This includes submodular, subadditive, coverage, and weighted-sum-of-matroid-rank valuations. However, two interesting counter-examples come to mind: multi-unit (where items are indistinguishable), and single-minded valuations. Nevertheless, the underlying optimization problem is not APX hard for multi-unit valuations, and for single-minded valuations the computational hardness of approximation is $1/k^{1-\epsilon}$ even without the extra constraint of truthfulness (see [4]).

Our second hardness result pertains to deterministic mechanisms for a very simple, non-regular class: budgeted additive valuations. This is despite the fact that the underlying k -player optimization problem admits an FPTAS [2]. Budgeted additive valuations are defined as follows.

Definition 2.2. *We say a valuation $v : 2^M \rightarrow \mathbb{R}^+$ is budgeted additive if there exists a constant $B \geq 0$ (the budget) such that $v(A) = \min(B, \sum_{i \in A} v(\{i\}))$.*

2.3 MIR, Randomized MIR, MIDR, and MIWR

Maximal in range (MIR) algorithms were introduced in [15] as a paradigm for designing truthful approximation mechanisms for computationally hard problems. An algorithm \mathcal{A} is maximal-in-range if it induces a maximal-in-range allocation rule when k and m are fixed.

Definition 2.3. *A k -bidder, m -item allocation rule f is maximal-in-range (MIR) if there exists a set of allocations $\mathcal{R} \subseteq \mathcal{X}([m], [k])$, such that $\forall v_1, \dots, v_k$ $f(v_1, \dots, v_k) \in \arg \max_{S \in \mathcal{R}} \sum_i v_i(S_i)$.*

A generalization of maximal-in-range that uses randomization sometimes yields better algorithm. An algorithm \mathcal{A} is *randomized maximal-in-range* if it induces a maximal-in-range allocation rule for every realization of its random coins. It is well known a randomized MIR algorithm can be combined with the VCG payment scheme to yield universally truthful mechanisms.

Dobzinski and Dughmi defined a generalization of randomized maximal-in-range algorithms in [7], termed maximal-in-distributional-range (MIDR). Here, each element of the range is a distribution over allocations. The resulting mechanism outputs the distribution in the range that maximizes the expected welfare, and charges VCG payments.

Definition 2.4. *f is maximal-in-distributional-range (MIDR) if there exists a set \mathcal{D} of distributions over allocations such that for all v_1, \dots, v_k , $f(v_1, \dots, v_k)$ is a distribution $D \in \mathcal{D}$ that maximizes the expected welfare of a random sample from D .*

MIDR algorithms were used in [7] to obtain a polynomial-time truthful-in-expectation FPTAS for multi-unit auctions, despite a lower bound of 2 on polynomial-time maximal-in-range algorithms. Moreover, they exhibited a variant of multi-unit auctions for which an MIDR FPTAS exists, yet no deterministic (or even universally truthful) polynomial time mechanism can attain an approximation ratio better than 2. Notably, the MIDR algorithms presented in [7] are of the following special form.

Definition 2.5. *f is maximal in weighted range (MIWR) if f is MIDR, and moreover each distribution D in the range of f is a weighted allocation: There is a pure allocation $S \in \mathcal{X}([m], [k])$ such that D outputs S with some probability, and the empty allocation otherwise.*

We denote a weighted allocation that outputs S with probability w by the pair (w, S) . When there is room for confusion, we use the term *pure allocation* to refer to an unweighted allocation.

Our first result will apply to all polynomial time MIWR mechanisms, and is the first such negative result. In fact, this result also applies to any randomization over MIWR mechanisms, a class we term *randomized MIWR* mechanisms. Randomized MIWR mechanisms include all universally-truthful randomized MIR mechanisms as a special case. Our second result will apply to only randomized maximal-in-range mechanisms, yet applies to a very restricted class of valuations, namely budgeted-additive valuations.

2.4 Some Complexity Theory

Broadly speaking, our proof involves constructing a reduction that transforms every instance of a k -player mechanism design problem into an instance of one of k other problems $\mathcal{P}_1, \dots, \mathcal{P}_k$, each of which individually is presumed to be computationally hard. The reduction has the property that input instances with a given number of items, m , are all transformed into inputs of the same problem \mathcal{P}_i , but instances with a different number of items may map to a different one of the k

problems. This raises difficulties because the complexity of $\mathcal{P}_1, \dots, \mathcal{P}_k$ may be “wild”: for each of them, there may be some input sizes (perhaps even infinitely many) that can be solved by a polynomial-sized Boolean circuit. In this section we develop the relevant complexity-theoretic tools to surmount this obstacle. We relegate the proofs of these results to Appendix C.

Definition 2.6. A set $S \subseteq \mathbb{N}$ is said to be complexity-defying (CD) if there exists a family of polynomial-sized Boolean circuits $\{C_n\}_{n \in \mathbb{N}}$ such that for all $n \in S$, the circuit C_n correctly decides 3SAT on all instances of size n .

A set $T \subseteq \mathbb{N}$ is said to be polynomially complexity-defying (PCD) if there exists a complexity-defying set S and a polynomial function $p(n)$ such that $T \subseteq \bigcup_{n \in S} [n, p(n)]$. Here $[a, b]$ denotes the set of all natural numbers x such that $a \leq x \leq b$. If a set $U \subseteq \mathbb{N}$ is not PCD, we say it is non-PCD.

Lemma 2.7. A finite union of CD sets is CD, and a finite union of PCD sets is PCD.

Definition 2.8. A decision problem or promise problem is said to have the padding property if for all $n < m$ there is a reduction that transforms instances of size n to instances of size m , running in time $\text{poly}(m)$ and mapping “yes” instances to “yes” instances and “no” instances to “no” instances. Similarly, an optimization problem is said to have the padding property if for all $n < m$ there is a reduction that transforms instances of size n to instances of size m , running in time $\text{poly}(m)$ and preserving the optimum value of the objective function.

Lemma 2.9. Suppose that \mathcal{L} is a decision problem or promise problem that has the padding property and is NP-hard under polynomial-time many-one reductions. Let T be any subset of \mathbb{N} . If there is a polynomial-sized circuit family that decides \mathcal{L} correctly whenever the input size belongs to T , then T is PCD.

Lemma 2.10. If \mathbb{N} is PCD, then $\text{NP} \subseteq \text{P}/\text{poly}$.

2.5 Technical Assumptions For Main Result

For our main result, a note is in order on the representation of valuation. Our results hold in the computational model. Therefore, we may assume that valuation functions are *succinct*, in that they are given as part of the input, and can be evaluated in time polynomial in the length of their description. Naturally, our main result applies to non-succinct valuations with oracle access, when the resulting problem admits a suitable reduction from an APX hard optimization problem.

Moreover, due to the generality of our main result, we need to make some technical assumptions. Namely, we restrict our attention to Combinatorial Auctions over a well-behaved family of instances. This restriction is without loss of generality for all well-studied classes of valuations for which the problem is APX-hard, such as coverage, submodular, etc. A family I of inputs to Combinatorial auctions is *well-behaved* if there exists a polynomial $b(m)$ such that for each input $(k, m, v_1, \dots, v_k) \in I$, the function v_i is represented as a bit-string of length $O(b(m))$, and moreover always evaluates to a rational number with $O(b(m))$ bits. While we believe this assumption may be removed, we justify it on two grounds: First, every well-studied variant of combinatorial auctions that is APX hard is also APX hard on a well-behaved family of instances, so this restriction is without loss for all such variants. Second, this assumption greatly simplifies our proof, since it allows us to describe the size of an instance by a single parameter, namely m .

3 Amplified Hardness for APX-Hard Valuations

In this section, we prove the following main result.

Theorem 3.1. *Fix a regular valuation class \mathcal{C} for which 2-player social welfare maximization is APX-hard. Fix a constant $k \geq 1$. For any constant $\epsilon > 0$, no polynomial-time randomized MIWR algorithm for k -player combinatorial auctions achieves an expected approximation ratio of $1/k + \epsilon$, unless $\text{NP} \subseteq \text{P/Poly}$.*

It is worth noting that this impossibility result applies to all universally-truthful randomized maximal-in-range algorithms. First, we prove the analogous result for MIWR mechanisms that take polynomial advice.

Theorem 3.2. *Fix a regular valuation class \mathcal{C} for which 2-player social welfare maximization is APX-hard. Fix a constant $k \geq 1$. For any constant $\epsilon > 0$, no non-uniform polynomial-time MIWR algorithm for k -player combinatorial auctions achieves an expected approximation ratio of $1/k + \epsilon$, unless $\text{NP} \subseteq \text{P/Poly}$.*

We then complete the proof by showing that any randomized MIWR mechanism can be “de-randomized” to one that takes polynomial advice.

A word is in order on the notion of *non-uniform* computation. For this, the reader should refer to Appendix D.3. Our hardness results in this section follow from the commonly-held conjecture that non-uniform computation cannot solve NP-complete problems, in other words $\text{NP} \not\subseteq \text{P/Poly}$.

Our proof strategy is as follows. In Section 3.1 we define a “perfect valuation profile” on k players as a set of valuations where exactly one player is interested in each item. We then show that any range of allocations that gives a good approximation on a randomly drawn perfect valuation profile must “shatter” a constant fraction of the items, meaning that the range contains all allocations of that subset of the items to q of the players, where the value of q depends on the quality of the approximation. (A better approximation implies a larger q .)

In Section 3.2, we prove Theorem 3.2 by induction on the number of players k . Roughly speaking, we show that for any MIWR mechanism \mathcal{A} for k players, the allocations with weight much larger than $1/k + \epsilon$ are useless. Namely, the inductive hypothesis implies that the allocations with weight sufficiently larger than $1/k + \epsilon$ cannot yield a good approximation to a randomly drawn perfect valuation; otherwise, one could use the resulting shattered set of items to design a strictly better MIWR mechanism for k' players for some $k' < k$. This allows us to conclude that all “useful” allocations have very similar weight to one another; within $1 - \eta$ for arbitrarily small η and a sufficiently large set of items. Since the mechanism maximizes over a large set of allocations that are almost “pure”, in the sense that the weights are almost identical, this yields a PTAS, contradicting the APX hardness of the problem.

Finally, we complete the proof of Theorem 3.1 in Section 3.3, using a de-randomization argument. This step is similar to Adleman’s proof that $\text{BPP} \subseteq \text{P/Poly}$.

3.1 Perfect Valuations

We define a perfect valuation profile as one where each item is desired by exactly one player. Perfect valuation profiles will prove useful in our proof, due to the fact that no “small” range can well-approximate social welfare for a randomly-drawn perfect valuation profile.

Definition 3.3. Let N and M be a set of players and items, respectively. Let $v_i : 2^M \rightarrow \mathbb{R}^+$ be the valuation of player $i \in N$. We say the valuation profile $\{v_i\}_{i \in N}$ is a perfect valuation profile on N and M if there exists a total allocation S of M to N such that $v_i(j) = 1$ if $j \in S_i$, and $v_i(j) = 0$ otherwise. In this case, we say that $\{v_i\}_{i \in N}$ is the perfect valuation profile generated by S .

To use perfect valuations in our proof, they must be allowable valuations. Indeed, it follows immediately from definition 2.1 that any regular class of valuations contains all perfect valuations.

The key property of perfect valuations is a reinterpretation of Lemma 1.3, and can be summarized as follows. If a range \mathcal{R} of allocations achieves a “good” approximation for many perfect valuations, then \mathcal{R} must include all allocations of a constant fraction of the items to some set of q players. Here, the number of players q depends on the quality of the approximation guaranteed by \mathcal{R} , with a better approximation yielding a larger q . The precise dependence of q on the quality of the approximation, as stated in Lemma 1.3, will prove key in Section 3.2.

3.2 Hardness for Non-Uniform MIWR Mechanisms

In this section, we prove Theorem 3.2, assuming $\text{NP} \not\subseteq \text{P/Poly}$. We fix the valuation class \mathcal{C} as in the statement of the theorem. Moreover, we fix $\eta > 0$ such that the 2-player social welfare maximization problem is APX-hard to approximate within $1 - \eta$. The proof proceeds by induction on k . We need the following strong inductive hypothesis:

IH(k). For any constant $\alpha > 1/k$ and set $T \subseteq \mathbb{N}$, if a non-uniform polynomial-time MIWR algorithm for the k -player problem achieves an α -approximation for m items whenever $m \in T$, then T is PCD.

In other words, the set of input lengths for which any particular such algorithm may achieve an α -approximation is PCD. (See Section 2.4 for the definition of PCD.) It is clear that establishing **IH(k)** for all $k \geq 1$ proves Theorem 3.2, since \mathbb{N} is not PCD. The base case of $k = 1$ is trivial. We now fix k , and assume **IH(q)** for all $q < k$.

Assume for a contradiction that **IH(k)** is violated for some α . Let $\alpha > 1/k$ be the supremum over all values of α violating it. Note that **IH(k - 1)** implies that $\alpha \in \left(\frac{1}{k}, \frac{1}{k-1}\right]$. To simplify the exposition, we assume the supremum is attained, and fix the algorithm \mathcal{A} (and corresponding family of polynomial advice strings) achieving an α -approximation for all $m \in \mathbb{F}$ where \mathbb{F} is not PCD. Our arguments can all be easily modified to hold when the supremum is not attained, by instantiating \mathcal{A} to achieve $(\alpha - \zeta)$ instead, where $\zeta > 0$ is as small as needed for the forthcoming proof. The proof then proceeds as follows. Letting \mathcal{D}^m denote the range of \mathcal{A} when the number of items is m , we partition \mathcal{D}^m into k sets \mathcal{D}_q^m ($2 \leq q \leq k + 1$) according to the weight of the allocation. We also assign every $m \in \mathbb{F}$ to one or more subsets T_q ($2 \leq q \leq k + 1$); the definition of T_q is quite technical, but roughly speaking $m \in T_q$ if the output of \mathcal{A} , when applied to a random perfect valuation profile, has probability at least $1/k$ of being in \mathcal{D}_q^m . As we said, $\mathbb{F} = \bigcup_{q=2}^{k+1} T_q$. However, we will prove that T_q is PCD for all q , hence by Lemma 2.7 their union \mathbb{F} is PCD. This contradicts our earlier assumption that \mathbb{F} is not PCD and completes the proof.

To prove that T_q is PCD, we distinguish three cases depending on the value of q . If $2 \leq q \leq k - 1$, then we prove that $m \in T_q$ implies that there is a non-uniform polynomial-time MIWR mechanism for the q -player problem that achieves an approximation ratio strictly better than $1/q$ when the number of items is $\lceil \sigma m \rceil$, for some constant $\sigma > 0$. By our induction hypothesis, the set of all such $\lceil \sigma m \rceil$ is a PCD set. If $q = k$, then we proceed similarly but working with the k -player

problem and proving an approximation ratio strictly better than α when the number of items is $\lceil \sigma m \rceil$; once again this implies that the set of all such $\lceil \sigma m \rceil$ is a PCD set, by our hypothesis on α . Finally, if $q = k + 1$ then we prove that there is a non-uniform polynomial-time algorithm achieving approximation ratio $1 - \eta$ for the two-player social welfare maximization problem, where η was chosen so that the problem APX hard to approximate within $1 - \eta$. By Lemma 2.9, this implies T_q is PCD.

Defining the partition of the range. Recall that an MIWR mechanism fixes a range of weighted allocations for each m . Let \mathcal{D}^m denote the range of \mathcal{A} when the number of items is m . Let $\mathcal{R}^m = \{S \in \mathcal{X}([m], [k]) : (w, S) \in \mathcal{D}^m \text{ for some } w\}$ be the corresponding set of pure allocations. For each allocation $S \in \mathcal{R}^m$, we use $w(S)$ to denote the weight of S in \mathcal{D}^m . We assume without loss of generality that there is a unique choice of $w(S)$, since allocations with greater weight are always preferred. When $m \in \mathbb{F}$, we may assume without loss of generality that $w(S) \geq \alpha$ for every $S \in \mathcal{R}^m$, since \mathcal{A} achieves an α approximation for m . We fix $\epsilon > 0$ such that $\alpha > 1/k + \epsilon$, and $\xi > 0$ such that $1/k + \epsilon/2 = (1 - \xi)^{-1} \cdot (1/k)$, and let $\delta = \epsilon\eta/5k$. We partition \mathcal{R}^m into *weight classes* as follows:

- $\mathcal{R}_q^m = \{S \in \mathcal{R}^m : \frac{1}{(1-\xi^2)^q} \leq w(S) < \frac{1}{(1-\xi^2)^{(q-1)}}\}$, for $2 \leq q \leq k - 1$.
- $\mathcal{R}_k^m = \{S \in \mathcal{R}^m : \frac{\alpha}{1-\delta} \leq w(S) < \frac{1}{(1-\xi^2)^{(k-1)}}\}$
- $\mathcal{R}_{k+1}^m = \{S \in \mathcal{R}^m : \alpha \leq w(S) < \frac{\alpha}{1-\delta}\}$

We partition \mathcal{D}^m similarly: $\mathcal{D}_q^m = \{(w(S), S) : S \in \mathcal{R}_q^m\}$ for $2 \leq q \leq k + 1$.

Consider now the set \mathcal{V}^m of perfect valuation profiles on $[k]$ and $[m'] = \{1, \dots, m/2\}$, extended to $[m]$ by zero-extension. For a given $v \in \mathcal{V}^m$ and $2 \leq q \leq k$, let us say that $v \in \mathcal{V}_q^m$ if the set \mathcal{R}_q^m contains an allocation S that achieves at least a $(1 + \xi)(q - 1)/k$ approximation to the social welfare maximizer. Finally, let us say that $v \in \mathcal{V}_{k+1}^m$ if v does not belong to \mathcal{V}_q^m for any $q < k + 1$. Notice that if $v \notin \mathcal{V}_q^m$ then the best approximation ratio achievable using an allocation in \mathcal{D}_q^m is at most

$$\frac{(1 + \xi)(q - 1)}{k} \cdot \frac{1}{(1 - \xi^2)^{(q - 1)}} = \frac{1}{(1 - \xi)k} = \frac{1}{k} + \frac{\epsilon}{2} < \alpha. \quad (1)$$

However, by our assumption that \mathcal{A} achieves an α -approximation for all valuation profiles with $m \in \mathbb{F}$, the range \mathcal{D}^m must contain an α -approximation to the social welfare maximizer. If $m \in \mathbb{F}$ and $v \in \mathcal{V}_{k+1}^m$, therefore, it follows that \mathcal{D}_{k+1}^m must contain an α -approximation to the social welfare maximizer.

By the pigeonhole principle, at least one q satisfies

$$|\mathcal{V}_q^m| \geq \frac{1}{k} \cdot k^{m'}. \quad (2)$$

Let T_q denote the set of all $m \in \mathbb{F}$ such that (2) holds. By the preceding discussion, we have $\mathbb{F} = \cup_{q=2}^{k+1} T_q$. We now proceed to prove that T_q is a PCD set for all q , completing the proof.

Cases 1 and 2: Large weight classes ($q \leq k$). To each allocation S of m items to k players, we may associate a function $f_S : [m'] \rightarrow [k] \cup \{*\}$, that maps each item $x \in [m']$ to the player who receives that item in S , or $*$ if the item is unallocated. Similarly, to each perfect valuation profile v on $[k]$ and $[m']$ we may associate a function $f_v : [m'] \rightarrow [k]$ that maps each item to the unique player who assigns a nonzero valuation to that item. Note that S achieves a c -approximation to the

social-welfare-maximizing allocation for v if and only if the functions f_S and f_v differ on $(1 - c)m'$ or fewer elements of $[m']$.

Assume now that $q \leq k$. If $m \in T_q$ then at least $1/k$ fraction of all perfect valuation profiles in \mathcal{V}_q^m have an allocation $S \in \mathcal{R}_q^m$ that achieves a $(1 + \xi)(q - 1)/k$ -approximation to the maximum social welfare. Thus, for at least $1/k$ fraction of all perfect valuation profiles $v \in \mathcal{V}_q^m$, there is some $S \in \mathcal{R}_q^m$ such that the f_S and f_v differ on $\left(1 - \frac{q-1}{k} - \frac{(q-1)\xi}{k}\right) m'$ elements of $[m']$. Applying Lemma 1.3, there is a set W of at least $\lceil \sigma m \rceil$ elements of $[m']$, and a set N' of q players in $[k]$, such that all allocations of W to N' occur as restrictions of allocations in \mathcal{R}_q^m . We refer to W as a “shattered” subset of $[m']$.

When $q < k$ (Case 1 of our argument) we may now construct, via a non-uniform polynomial-time reduction, an MIWR allocation rule for the q -player problem that achieves a $[(1 - \xi^2)q]^{-1}$ approximation for $\lceil \sigma m \rceil$ items when $m \in T_q$. Using W and N' – as defined above – as advice, embed the instance into an input for \mathcal{A} by using players N' and items W in the obvious way: Give player in $[k] \setminus N'$ an all-zero valuation. Moreover, extend the valuation of a player $i \in N'$ to the entire set of items $[m]$. Now, run \mathcal{A} on the embedded instance. Notice that every allocation of W to N' appears as the restriction of some allocation in \mathcal{R}_q^m , and is therefore in the range of \mathcal{A} with weight at least $[(1 - \xi^2)q]^{-1}$. Thus, \mathcal{A} must output a weighted allocation with expected welfare at least $[(1 - \xi^2)q]^{-1}$ of the optimal. The result is a non-uniform poly-time MIWR mechanism for q players with approximation ratio bounded away from $1/q$ for all integers $\hat{m} = \lceil \sigma m \rceil$ such that $m \in T_q$. By our induction hypothesis, this implies that the sum of all such \hat{m} is a PCD set. The fact that T_q itself is a PCD set now follows as an easy application of the definition of PCD.

When $q = k$ (Case 2 of our argument) using the same embedding yields an algorithm for k players that achieves an $\alpha/(1 - \delta)$ approximation for all $\hat{m} = \lceil \sigma m \rceil$ such that $m \in T_q$. By our definition of α , this implies that the set of all such \hat{m} is a PCD set, which again implies that T_q is a PCD set.

Case 3: The smallest weight class ($q = k + 1$). The remaining case is $q = k + 1$. When $m \in T_{k+1}$, by our definition of \mathcal{V}_{k+1}^m , at least $1/k$ fraction of all (extended) perfect valuation profiles $v \in \mathcal{V}^m$ have a weighted allocation $(w(S), S) \in \mathcal{D}_{k+1}^m$ that is an α -approximation to the social welfare maximizing allocation for v . Since $\alpha \leq w(S) < \alpha/(1 - \delta)$, the pure allocation S must be a $(1 - \delta)$ -approximation to the social welfare maximizer. On the other hand, our assumption is that maximizing social welfare is APX-hard, even for two players; to be specific, recall that $\eta > 0$ was chosen such that it is NP-hard to approximate the maximum social welfare with approximation factor $1 - \eta$. We now complete the proof by exhibiting a randomized, non-uniform polynomial time algorithm that achieves a $(1 - \eta)$ -approximation for the k -player problem with $m/2$ items, for all $m \in T_{k+1}$. Notice that the de-randomization argument of Adleman [1] for proving $\text{BPP} \subseteq \text{P/Poly}$ can be used to de-randomize this to a non-uniform deterministic $(1 - \eta)$ -approximation for the k -player problem with $m/2$ items, for all $m \in T_{k+1}$. The reader unfamiliar with Adleman’s argument may refer to Section 3.3, where we use the argument to establish Theorem 3.1.

Recall that $\delta = \epsilon\eta/5k$. We will now use \mathcal{A} to get a $(1 - \eta)$ -approximate solution for an instance with k players and $m' = m/2$ items for all $m \in T_{k+1}$. We embed the instance on k players and m' items into \mathcal{A} in the following way. We use $M_1 = [m] \setminus [m']$ and let $v_i : 2^{M_1} \rightarrow \mathbb{R}$ denote the resulting valuation of player i . We assume without loss of generality that $\max_i v_i(M_1) = 1$. Next, we modify each player’s valuation function by “mixing in” a perfect valuation profile on the remaining set of items $M_2 = [m']$. We draw a perfect valuation profile (v'_1, \dots, v'_k) on N and M_2 uniformly at

random. Now, we “mix” the original valuations v with v' , in proportions 1 and $\gamma = \frac{4k}{\epsilon m}$, to yield the following *hybrid valuation profile* $v^* : 2^M \rightarrow \mathbb{R}^+$.

$$v_i^* = v_i \oplus \gamma v'_i$$

We abuse notation and use $v_i [v'_i]$ to refer also to the zero-extension of $v_i [v'_i]$ to M . Let $OPT = \max_{S \in \mathcal{X}} v(S)$. Similarly, let $OPT' = \max_{S \in \mathcal{X}} v'(S)$ and let $OPT^* = \max_{S \in \mathcal{X}} v^*(S)$. Notice that $1 \leq OPT \leq k$, and that $OPT' = m'$, by construction. Since v and v' are defined on a disjoint set of items, it is easy to see that $OPT^* = OPT + \gamma OPT'$. The scalar γ was carefully chosen so that the following facts hold:

1. The random valuation profile v' accounts for a majority share of v^* in any optimal solution. Specifically, $\gamma OPT' \geq \frac{4}{\epsilon} OPT$. This implies that an approximation to the optimal welfare using v^* gives a similar approximation to the optimal welfare using v' . To be more precise, it can be shown by a simple calculation that:

Claim 3.4. *For any $S \in \mathcal{X}$ and any $\beta \geq 0$, if $v^*(S) \geq \beta OPT^*$ then $v'(S) \geq (\beta - \epsilon/2) OPT'$.*

2. The original valuation profile v accounts for a constant-factor share of v^* in any optimal solution. Specifically $OPT \geq \frac{\epsilon}{4k} (\gamma OPT')$. This implies that a $(1 - \delta)$ -approximation to the optimal welfare using v^* gives a $(1 - O(\delta))$ -approximation to the optimal welfare using v . To be more precise, it can be shown by a simple calculation that:

Claim 3.5. *For any $S \in \mathcal{X}$, if $v^*(S) \geq (1 - \delta) OPT^*$ then $v(S) \geq (1 - \frac{5k}{\epsilon} \delta) OPT = (1 - \eta) OPT$.*

We are now ready to show that running \mathcal{A} on the valuations v^* will yield, with constant probability, an allocation that is a $(1 - \eta)$ -approximation to the optimal welfare for the original valuations v , when $m \in T_{k+1}$. Let $(w(S), S)$ be the weighted allocation output by \mathcal{A} ; note that S is a random variable over draws of v' . Since \mathcal{A} is an α approximation algorithm, the welfare $w(S)v^*(S)$ is at least $\alpha OPT^* \geq (1/k + \epsilon) OPT^*$ with probability 1. This implies that $v^*(S) \geq \left(\frac{1}{w(S) \cdot k} + \frac{\epsilon}{w(S)} \right) OPT^*$. By Claim 3.4, we see that $v'(S)$ is not too far behind: $v'(S) \geq \left(\frac{1}{w(S) \cdot k} + \frac{\epsilon}{w(S)} - \frac{\epsilon}{2} \right) OPT'$. Moreover, this gives:

$$w(S)v'(S) \geq \left(\frac{1}{k} + \frac{\epsilon}{2} \right) OPT' \tag{3}$$

Recall from equation (1) that if $v' \in \mathcal{V}_{k+1}^m$ then for $2 \leq q \leq k$, there is no $S \in \mathcal{R}_q^m$ that satisfies (3), hence any such S satisfying (3) must belong to \mathcal{D}_{k+1}^m . Also, by our assumption that $m \in T_{k+1}$, the probability that $v' \in \mathcal{V}_{k+1}^m$ is at least $1/k$.

We have thus established that running \mathcal{A} on the random input v^* yields, with probability at least $1/k$, an outcome $(w(S), S)$ in \mathcal{D}_{k+1}^m . Using the fact that $w \leq \alpha/(1 - \delta)$ and $w(S)v^*(S) \geq \alpha OPT$, we conclude that S is $(1 - \delta)$ -approximate for v^* also with probability $1/k$:

$$v^*(S) \geq (1 - \delta) OPT^*$$

Invoking Claim 3.5, we conclude that $v(S) \geq (1 - \eta) OPT$ with constant probability over draws of v' . Since $w(S)$ is at least $1/k$, S is output by \mathcal{A} with constant probability. This completes the proof.

3.3 Main Result

In this section, we complete the proof of Theorem 3.1. First, we make the observation that running a randomized MIWR algorithm multiple times independently and returning the best allocation output by any of the runs results in another randomized MIWR algorithm.

Lemma 3.6. *Fix a randomized MIWR algorithm \mathcal{A} and a positive integer r . Let \mathcal{A}^r be the algorithm that runs r independent executions of \mathcal{A} on its input, and of the r allocations returned, outputs the one with greatest welfare. \mathcal{A}^r is also randomized MIWR.*

Proof. Condition on $\mathcal{D}_1, \dots, \mathcal{D}_r$, the ranges of \mathcal{A} on the r independent executions. \mathcal{A} maximizes expected welfare over \mathcal{D}_i on execution i . Therefore \mathcal{A}^r maximizes over $\mathcal{D}_1 \cup \dots \cup \mathcal{D}_r$. \square

Now, we derive Theorem 3.1 from Theorem 3.2, using a de-randomization argument similar to that of Adleman [1]. Assume for a contradiction that \mathcal{A} is a randomized MIWR algorithm that runs in polynomial time and achieves an expected approximation ratio $1/k + \epsilon$ for each input m and v_1, \dots, v_k . Let n denote the number of bits in the input, and let $\ell(n)$ be a polynomial bounding the length of the random string drawn by \mathcal{A} . We will describe a polynomial-time with polynomial-advice MIWR algorithm that achieves an approximation ratio of $1/k + \epsilon/2$, which contradicts Theorem 3.2.

Let $r(n) = 2n/\epsilon^2$ and let $\mathcal{A}' = \mathcal{A}^{r(n)}$. By Lemma 3.6, \mathcal{A}' is randomized MIWR, runs in polynomial time, and draws at most $\ell(n)r(n)$ random bits. Let X_i be the fraction of the optimal social welfare achieved by the allocation output on the i 'th run of \mathcal{A} . The random variables $X_1, \dots, X_{r(n)}$ are independent, $0 \leq X_i \leq 1$, and $E[X_i] \geq 1/k + \epsilon$. For each input of length n , the probability that none of the $r(n)$ runs of \mathcal{A} return an allocation with welfare better than $1/k + \epsilon/2$ of the optimal can be upper-bounded using Hoeffding's inequality:

$$Pr \left[\max_i X_i \leq \left(\frac{1}{k} + \frac{\epsilon}{2} \right) \right] \leq Pr \left[E \left(\sum_i X_i \right) - \sum_i X_i \geq \frac{\epsilon r(n)}{2} \right] \leq e^{-\epsilon^2 r(n)/2} = e^{-n}.$$

The number of different inputs of length n is 2^n . Thus, using the union bound and the above inequality, the probability that \mathcal{A} outputs a $(1/k + \epsilon/2)$ -approximate allocation on all inputs of length n is non-zero. Therefore, for each n there is choice of at most $\ell(n)r(n)$ random bits such that \mathcal{A}' achieves a $1/k + \epsilon/2$ approximation for all inputs. Using this as the advice string, this contradicts Theorem 3.2. This completes the proof of Theorem 3.1.

4 Hardness Result for Budgeted Additive Valuations

In this section, we prove the following theorem:

Theorem 4.1. *There is no polynomial time randomized MIR mechanism that achieves $\frac{1}{k} + \epsilon$ approximation of the optimal social welfare for k bidders with budgeted additive valuations, unless $\text{NP} \subseteq \text{P/Poly}$.*

Notice that this theorem implies that all universally truthful randomized MIR mechanisms cannot achieve $\frac{1}{k} + \epsilon$ approximation for k bidders. In the proof we will use the term k -partition interchangeably with an allocation for k bidders. As in the previous section, an allocation does not

necessarily allocate all items. A partition corresponding to a total allocation will be called *covering*. Perfect valuations generated by total allocations will also be used in the proof. We first study the abundance of “orthogonal” partitions of $M = [m]$. The following definition formalizes this notion.

Definition 4.2. Let \mathcal{T} be a set of k -partitions of M : $\mathcal{T} = \{(T_1^i, T_2^i, \dots, T_k^i) \mid i \in [\ell]\}$, we say these partitions are ϵ -apart for $\epsilon > 0$, if,

$$\forall (I_1, I_2, \dots, I_\ell) \in [k]^\ell, \quad \left| \bigcap_{i=1}^{\ell} T_{I_i}^i \right| \leq \left(\frac{1}{k} \right)^\ell (1 + \epsilon) |M|.$$

Lemma 4.3. For every pair of integers k and ℓ , and every ϵ satisfying $0 < \epsilon < 1$, there exists $\alpha > 0$ such that there exists a set F of covering k -partitions of M , where $|F| = e^{\alpha m}$, and every ℓ elements of F are ϵ -apart.

The proof of Lemma 4.3 is relegated to the appendix. The next lemma shows that for valuations generated by partitions that are apart, good welfare approximations require distinct allocations.

Lemma 4.4. Given ℓ covering k -partitions that are ϵ -apart, consider the ℓ tuples of valuations generated by them. The sum of social welfare achievable by a single allocation on these tuples of valuations is at most $\left(\frac{1}{k} + \sqrt{\frac{k\pi}{\ell}} + \frac{k2^{-\ell/k}}{\ell \ln 2} \right) (1 + \epsilon)$ of the optimal.

Proof. The sum of optimal social welfare for ℓ generated valuations is easily seen to be ℓm . Let $\mathcal{T} = \{(T_1^i, T_2^i, \dots, T_k^i) \mid i \in [\ell]\}$ be the set of covering k -partitions that generate the valuations. Let S be the sum of social welfare achieved by one single allocation R , then S is maximized when $R(x) \in \operatorname{argmax}_{i \in [k]} \sum_{j=1}^{\ell} v_i^j(x)$ holds for every item $x \in M$. For $I \in [k]^\ell$, define $Q_i(I)$ to be the number of times that i occurs in I , i.e., $|\{j \in [\ell] \mid I_j = i\}|$, and define the *plenty* of I to be $P(I) = \max_{i \in [k]} Q_i(I)$. Then we have

$$S \leq \sum_{I \in [k]^\ell} P(I) \cdot |T_{I_1}^1 \cap T_{I_2}^2 \cap \dots \cap T_{I_\ell}^\ell| \leq \frac{\sum_{I \in [k]^\ell} P(I)}{k^\ell} (1 + \epsilon) m. \quad (4)$$

The second inequality results from the ϵ -apartness of the partitions.

We recognize that $\frac{\sum_{I \in [k]^\ell} P(I)}{k^\ell}$ can be seen as the expectation of a properly defined random variable — if I' is a random variable uniformly distributed on $[k]^\ell$, then this factor is exactly the expectation of $P(I')$. The problem boils down to bounding $E[P(I')]$. Note that $E[Q_i(I')] = \ell/k$. Let $Y(I) = P(I) - \frac{\ell}{k}$, then by the union bound

$$\Pr[Y(I') > \delta \cdot \frac{\ell}{k}] \leq k \cdot \Pr[Q_i(I') > (1 + \delta) \frac{\ell}{k}].$$

Applying Chernoff bound (Theorem B.1), we get

$$\begin{aligned} E[Y(I')] &= \int_0^\infty \Pr_{I'}[Y(I') > \delta] d\delta = \frac{\ell}{k} \int_0^\infty \Pr_{I'}[Y(I') > \left(\delta \cdot \frac{k}{\ell} \right) \frac{\ell}{k}] d\left(\delta \cdot \frac{k}{\ell} \right) \\ &\leq \ell \left\{ \int_0^{2e-1} \Pr_{I'}[Q_i(I') > (1 + \delta) \frac{\ell}{k}] d\delta + \int_{2e-1}^\infty \Pr_{I'}[Q_i(I') > (1 + \delta) \frac{\ell}{k}] d\delta \right\} \\ &< \ell \left(\int_0^\infty e^{-\delta^2 \ell / 4k} d\delta + \int_1^\infty 2^{-\delta \ell / k} d\delta \right) = \sqrt{k\pi\ell} + \frac{k2^{-\ell/k}}{\ln 2}. \end{aligned}$$

Consequently we also obtain that $E[P(I')] \leq \frac{\ell}{k} + \sqrt{k\pi\ell} + \frac{k2^{-\ell/k}}{\ln 2}$. Substituting this into (4), then dividing it by ℓm , the sum of optimal social welfare, we get what the lemma claims. \square

To finish the proof, we need the next lemma that connects to the “shattering” lemma.

Definition 4.5. *Two k -partitions (T_1, T_2, \dots, T_k) and $(T'_1, T'_2, \dots, T'_k)$ are said to be ϵ -far if $\sum_{i \neq j} |T_i \cap T'_j| \geq \epsilon m$. If two partitions are not ϵ -far, we say that they are ϵ -close.*

Lemma 4.6. *If an MIR mechanism achieves at least $\frac{1}{k} + \sqrt{\frac{k\pi}{\ell}} + \frac{ke^{-\ell/k}}{\ell \ln 2} + \epsilon'$ approximation to the optimal social welfare for k bidders, where k , ℓ and ϵ' are all fixed, and $\ell \geq 10k$, then there is a $\delta > 0$, such that there is a subset S of items, with $|S| \geq \delta m$, and two bidders i and j , and every allocation of items in S to i and j is a restriction of an allocation in the range of the mechanism.*

Proof. Let ϵ be $\epsilon'/3$, and by Lemma 4.3 there is a set F of ϵ -apart k -partitions, and $|F| = e^{\alpha|U|}$ for some $\alpha > 0$. Let \mathcal{R} be the multi-set of allocations output by the mechanism on the tuples of valuations generated by the partitions in F . Note that $|\mathcal{R}| = |F|$. We claim that in \mathcal{R} , there can be no ℓ partitions such that every two of them are $\frac{\epsilon}{\ell^2}$ -close. For a contradiction, suppose that this is the case. Let $\{(T_1^i, T_2^i, \dots, T_k^i) \mid i \in [\ell]\}$ be these allocations, define D to be $\bigcup_{1 \leq s < t \leq \ell} \bigcup_{i \neq j} (T_i^s \cap T_j^t)$, then because of the pairwise $\frac{\epsilon}{\ell^2}$ -closeness of the partitions, $|D| \leq \epsilon m$. For each item not in D , the ℓ allocations either allocate it in the same way, or some allocate it in the same way and others do not allocate it to any bidder. By Lemma 4.4, on $M \setminus D$ the allocations can achieve at most $\frac{1}{k} + \sqrt{\frac{k\pi}{\ell}} + \frac{k2^{-\ell/k}}{\ell \ln 2} + \epsilon$ of the sum of optimal welfare. Each item in D can contribute at most ℓ to the sum of welfare, and in total they count at most ϵ fraction of the optimal, which is ℓm . Thus the mechanism can achieve at most $\frac{1}{k} + \sqrt{\frac{k\pi}{\ell}} + \frac{k2^{-\ell/k}}{\ell \ln 2} + \frac{2\epsilon'}{3}$ of the optimal social welfare, contradicting the assumption on its performance. Thus, there are no ℓ -allocations in \mathcal{R} that are pairwise $\frac{\epsilon}{\ell^2}$ -far. Applying Lemma 1.4 to \mathcal{R} , we finish the proof. \square

Proof of Theorem 4.1: Whenever we have the range of a mechanism containing all allocations of items in a linearly smaller subset to two bidders, we can use the mechanism with polynomial advice to optimize the social welfare of an auction with fewer items. Therefore, the condition of Lemma 4.6 should not be satisfied for any ℓ unless $\text{NP} \subseteq \text{P/Poly}$. Let ℓ in Lemma 4.6 get arbitrarily big. We see that, unless $\text{NP} \subseteq \text{P/Poly}$, any efficient MIR mechanism cannot achieve $1/k + \epsilon$ approximation to the optimal social welfare for k bidders. Then by the same argument as in Section 3.3 (proof omitted here), we can extend this to randomized MIR mechanisms and get Theorem 4.1. \square

Acknowledgements

We thank the authors of [5] and [13] — Elchanan Mossel, Christos Papadimitriou, Michael Schapira, Yaron Singer, Dave Buchfuhrer, and Chris Umans — for stimulating discussions on these topics and for their influence on this work.

References

- [1] Leonard Adleman. Two theorems on random polynomial time. In *SFCS '78: Proceedings of the 19th Annual Symposium on Foundations of Computer Science*, pages 75–83, Washington, DC, USA, 1978. IEEE Computer Society.
- [2] Nir Andelman and Yishay Mansour. Auctions with budget constraints. In *SWAT*, pages 26–38, 2004.
- [3] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 1 edition, April 2009.
- [4] Liad Blumrosen and Noam Nisan. 2007. Combinatorial Auctions (a survey). In “Algorithmic Game Theory”, N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani, editors.
- [5] Dave Buchfuhrer and Chris Umans. Limits on the social welfare of maximal-in-range auction mechanisms. working paper.
- [6] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.
- [7] Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. In *FOCS'09*. to appear.
- [8] Shahar Dobzinski and Noam Nisan. Limitations of VCG-based mechanisms. Preliminary version in STOC'07.
- [9] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. In EC'07.
- [10] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In STOC'06.
- [11] T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.
- [12] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In FOCS 2005.
- [13] Elchanan Mossel, Christos Papadimitriou, Michael Schapira, and Yaron Singer. VC v. VCG: Inapproximability of combinatorial auctions via generalizations of the VC dimension. Working Paper.
- [14] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in STOC 1999.
- [15] Noam Nisan and Amir Ronen. Computationally feasible VCG-based mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007. A preliminary version appeared in EC 2000.
- [16] Christos Papadimitriou, Michael Schapira, and Yaron Singer. On the hardness of being truthful. In *FOCS*, 2008.
- [17] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

A Shattering Results

We first formally define the notion of “shattering” in a more general setting.

Definition A.1. For any sets U, V we interpret the notation V^U to mean the set of functions from U to V . If $R \subseteq V^U$, $S \subseteq U$, $L \subseteq V$, we say that S is (L, q) -shattered by R , for an integer q , $2 \leq q \leq |L|$, if there exist q functions $c_1, c_2, \dots, c_q : S \rightarrow L$ that satisfy:

1. $\forall x \in S \quad \forall i \neq j \quad c_i(x) \neq c_j(x)$
2. $\forall h \in [q]^S \quad \exists f \in R \quad \forall x \in S \quad f(x) = c_{h(x)}(x)$

Intuitively, we associate with each element in S a range in L of size exactly q , and we say that S is (L, q) -shattered by R if every function that maps each element in S to its associated range is a restriction of an element in R . In the context of combinatorial auctions, we see U as the set of items, and V as the set of bidders, plus a dummy bidder representing not allocating the item. Then set of functions V^U is the set of all possible allocations.

The following observation bridges this notion of shattering to its application to the combinatorial auctions in the paper.

Observation A.2. If a subset S of size δm is (L, q) -shattered by $R \subseteq V^U$, then there exists a subset $L' \subseteq L$ and $S' \subseteq S$, such that $|L'| = q$, $|S'| \geq |S| / \binom{|L|}{q}$ and S' is (L', q) -shattered by R .

The observation is easily seen by the pigeonhole principle. Note that by the definition of (L, q) -shattering, if $|L'| = q$, then we have that every function from S' to L' is a restriction of an element in R . In the context of combinatorial auctions, this means that all possible allocations of items in S' to the q bidders in L' are in the range R under restriction. It is this form of “strong” shattering that is in use in the main body of the paper. In the following lemmas, we will show the existence of large (L, q) -shattered sets, being aware that an application of the above observation implies a subset being “strongly” shattered, of size only a constant factor smaller.

Lemma A.3. For every integer $n \geq 2$, q , $2 \leq q \leq n$, and every $\epsilon > 0$, there is a $\delta > 0$ such that the following holds. For every pair of finite sets M, N with $|N| = n$ and every set R of more than $(q - 1 + \epsilon)^{|M|}$ elements of N^M there is a set S of at least $\delta |M|$ elements of M such that S is (V, q) -shattered by R .

Proof. Let $F_q(m, n, d)$ denote the maximum cardinality of a set $R \subseteq A^B$ such that $|A| = n$, $|B| = m$, and R does not (A, q) -shatter any $(d + 1)$ -element subset of B .

Fix an element $b \in B$. For each element $f \in R$, let f_{-b} denote the restriction of f to the set $B \setminus \{b\}$. Take the set of all functions $g : B \setminus \{b\} \rightarrow A$ and partition it into sets $Q_0, Q_1, \dots, Q_{\binom{n}{q}}$ as follows. First, given an ordered pair (g, a) consisting of a function g from $B \setminus \{b\}$ to A and an element $a \in A$, let $g * a$ denote the unique function f from B to A that maps b to a and restricts to g on $B \setminus \{b\}$. Now define $S(g)$ to be the set of all $a \in A$ such that $g * a$ is in R . Number all the q -element subsets of A from 1 to $\binom{n}{q}$, call them $P_1, P_2, \dots, P_{\binom{n}{q}}$, and let Q_i ($1 \leq i \leq \binom{n}{q}$) consist of all g such that $S(g)$ has at least q elements, and the q smallest elements of $S(g)$ constitute P_i . Finally let Q_0 consist of all g such that $S(g)$ has fewer than q elements.

By our assumption that R does not (A, q) -shatter any set of size greater than d , we have the following facts:

1. Q_0 does not (A, q) -shatter any $(d + 1)$ -element subset of $B \setminus \{b\}$. Consequently,

$$|Q_0| \leq F(m - 1, n, d).$$

2. For all $i \leq \binom{n}{q}$, Q_i does not (A, q) -shatter any d -element subset of $B \setminus \{b\}$. Consequently,

$$|Q_i| \leq F_q(m - 1, n, d - 1).$$

Let R_i denote the set of all $f \in R$ such that f_{-b} is in Q_i , for $0 \leq i \leq \binom{n}{q}$, then by definition of Q_i , we have $|R_0| \leq (q - 1)|Q_0|$, and $|R_i| \leq n|Q_i|$ for $i \leq 1$. Since R_i 's are disjoint, we have

$$|R| = \sum_{i=0}^{\binom{n}{q}} |R_i| \leq (q - 1)|Q_0| + \sum_{i=1}^{\binom{n}{q}} n|Q_i|,$$

$$F_q(m, n, d) \leq (q - 1)F_q(m - 1, n, d) + n \binom{n}{q} F_q(m - 1, n, d - 1) \quad (5)$$

The recurrence (5), together with the initial condition $F_q(m, n, 0) = (q - 1)^m$ for all m, n , implies the upper bound

$$F_q(m, n, d) \leq \sum_{i=0}^d n^i \binom{n}{q}^i \binom{m}{i} (q - 1)^m$$

Thus, if $F_q(m, n, d) > (q - 1 + \epsilon)^m$ then, by using Stirling's approximation, we see that $d > \delta m$ for some δ depending only on ϵ and n . \square

In Section 3 of the paper, we made use of the fact that a range of allocations shatters a large subset if they generate good social welfare for many perfect valuations. The condition is captured by the following definition:

Definition A.4. For two functions $f, g \in N^M$, their normalized Hamming distance $\text{Ham}(f, g)$ is equal to $\frac{1}{|M|}$ times the number of distinct $x \in M$ such that $f(x) \neq g(x)$. If $f \in N^M$ and $R \subseteq N^M$, the Hamming distance $\text{Ham}(f, R)$ is the minimum of $\text{Ham}(f, g)$ for all $g \in R$.

As each perfect valuation can be seen as a function f in N^M , and each allocation can be viewed as a $g \in N^M$, $\text{Ham}(f, g)$ is how much social welfare is lost by g on the perfect valuation f . In the same way, R can be viewed as a range of allocations, and $\text{Ham}(f, R)$ is the minimum social welfare lost by any of the allocation in R on valuation f . If $\text{Ham}(f, R)$ is small for a large fraction of $f \in N^M$, it means the range achieves a good approximation of social welfare for a significant portion of the perfect valuations.

We also note that since N can represent the set of bidders plus a dummy bidder representing not allocating an item, N^M can express all allocations including those not allocating all items. On the other hand, if we restrict the functions so that they can take values only in a subset L representing the real bidders, then they represent allocations that do not discard items. This explains the role played by the set L in the next lemma.

Lemma A.5. For every positive real number $\epsilon > 0$, integers $n \geq 2$, q , $2 \leq q \leq n$, and polynomial $\gamma(n)$, there is a $\delta > 0$ such that the following holds. For all finite sets M, N and all subsets $L \subseteq N$ with $|L| = n$, if $R \subseteq N^M$ and at least $\gamma n^{|U|}$ points $f \in L^U$ satisfy $\text{Ham}(f, R) < 1 - (q - 1)/n - \epsilon$, then there is a set $S \subseteq M$ such that $|S| > \delta |M|$ and S is (L, q) -shattered by R .

Proof. Let $m = |M|, r = 1 - (q - 1)/n - \epsilon$. Let A be the set of all points $f \in L^M$ such that $\text{Ham}(f, R) < r$. Let G be a function from A to R such that $\text{Ham}(f, G(f)) < r$ for all $f \in A$. Let $I(f)$ denote the set of all $x \in M$ such that $f(x) = G(f)(x)$. Note that our assumption that $\text{Ham}(f, G(f)) < r$ implies that $|I(f)| \geq (\frac{q-1}{n} + \epsilon)m$. The number of pairs (f, J) such that $f \in A, |J| = \epsilon m/2, J \subseteq I(f)$ is bounded below by $\gamma n^m \cdot \binom{(1/n+\epsilon)m}{\epsilon m/2}$. By the pigeonhole principle, there is at least one set J of $\epsilon m/2$ elements such that the number of $f \in L^U$ satisfying $J \subseteq I(f)$ is at least

$$\begin{aligned} \gamma n^m \cdot \binom{(\frac{q-1}{n} + \epsilon)m}{\epsilon m/2} &/ \binom{m}{\epsilon m/2} = \gamma n^m \frac{((\frac{q-1}{n} + \epsilon)m)! ((1 - \epsilon/2)m)!}{((\frac{q-1}{n} + \epsilon/2)m)! m!} \\ &> \gamma n^m \cdot \frac{(\frac{q-1}{n} + \epsilon)m}{m} \cdot \frac{(\frac{q-1}{n} + \epsilon)m - 1}{m - 1} \cdots \frac{(\frac{q-1}{n} + \epsilon/2)m}{(1 - \epsilon/2)m} \\ &> \gamma n^m \left(\frac{\frac{q-1}{n} + \epsilon/2}{1 - \epsilon/2} \right)^{\epsilon m/2}. \end{aligned}$$

Fix such a set J . For every $f \in L^M$ satisfying $J \subseteq I(f)$, the restriction of f to J is an element $g \in L^J$; note that g is also the restriction of $G(f)$ to J . For any single $g \in L^J$, the number of $f \in L^M$ that restrict to g is bounded above by $n^{m-\epsilon m/2}$. Applying the pigeonhole principle again, we see that the number of distinct $g \in L^J$ that occur as the restriction of some $f \in A$ satisfying $J \subseteq I(f)$ must be at least

$$\gamma n^m \left(\frac{\frac{q-1}{n} + \epsilon/2}{1 - \epsilon/2} \right)^{\epsilon m/2} / n^{m-\epsilon m/2} = \gamma \left(\frac{q - 1 + \epsilon n/2}{1 - \epsilon/2} \right)^{\epsilon m/2}.$$

We now have the following situation. There is a set J of $\epsilon m/2$ elements, and at least $\gamma \cdot (q-1+\epsilon n/2)^{|J|}$ elements of L^J occur as the restriction of an element of R to J . It follows from Lemma A.3 that J has a subset of S of at least δm elements such that S is (L, q) -shattered by R . \square

Proof of Lemma 1.3: Combining Lemma A.3, Lemma A.5 and Observation A.2, we immediately get Lemma 1.3. \square

Now we are ready to show the next shattering lemma. The k -partitions represent allocations to k -bidders, allowed to discard items. In the proof we occasionally see them as partial functions, in a way very similar to that in the previous lemmas.

To state the lemma succinctly, we denote by $t(k, \ell, \epsilon, m, r)$ the smallest number of subsets that are $([k], 2)$ -shattered by any set \mathcal{T} of k -partitions of a set M , where $|M| = m, |\mathcal{T}| = r$, and every ℓ partitions from \mathcal{T} are not pairwise ϵ -close (see Definition 4.5).

Lemma A.6. *For every integers $k \geq 2, \ell$ and m , every $\epsilon > 0$, there exists an $\alpha > 0$ such that $t(k, \ell, \epsilon, m, r) \geq r^\alpha$ for every r .*

Note that if r is $2^{\beta m}$ for some $\beta > 0$, the conclusion implies the existence of a subset of size γm that is $([k], 2)$ -shattered by the set of partitions, for some $\gamma > 0$. The proof also works if R is a multi-set.

Proof. Let \mathcal{T} be any set of k -partitions of M , $|M| = m$, $|\mathcal{T}| = p$, and every ℓ elements from \mathcal{T} are not pairwise ϵ -close. We arbitrarily group the partitions in \mathcal{T} , so that every group consists of ℓ partitions. Then in each group $\{(T_1^i, T_2^i, \dots, T_k^i) \mid i \in [\ell]\}$, there are at least two partitions that are ϵ -far, and the size of their “difference” $\sum_{i \neq j} |T_i^s \cap T_j^t|$ is at least ϵm . Since we have r/ℓ such pairs, the sum of the sizes of “differences” will be at least $\frac{\epsilon m r}{\ell}$. By pigeonhole principle, there exists an $x \in M$, and $i^*, j^* \in [k]$ ($i^* \neq j^*$) such that in at least $\frac{\epsilon r}{\ell \binom{k}{2}}$ pairs of partitions (T_1, T_2, \dots, T_k) and $(T_1', T_2', \dots, T_k')$, x occurs in $(T_{i^*} \cap T_{j^*}') \cup (T_{i^*}' \cap T_{j^*})$. Now if we denote by \mathcal{T}_{i^*} the set of those partitions in \mathcal{T} that map x to i^* , and \mathcal{T}_{j^*} those mapping x to j^* , then $|\mathcal{T}_{i^*}| \geq \frac{\epsilon p}{\ell \binom{k}{2}}$ and $|\mathcal{T}_{j^*}| \geq \frac{\epsilon p}{\ell \binom{k}{2}}$.

Let I_{i^*} denote the set of subsets that are $([k], 2)$ -shattered by \mathcal{T}_{i^*} , and similarly I_{j^*} the set of subsets $([k], 2)$ -shattered by \mathcal{T}_{j^*} , I the set of subsets $([k], 2)$ -shattered by \mathcal{T} itself. We claim that $|I| \geq |I_{i^*}| + |I_{j^*}|$. To see this, it is clear that $I_{i^*} \cup I_{j^*} \subseteq I$, and x is not in any of the set in $I_{i^*} \cup I_{j^*}$. Besides, for every set S in $I_{i^*} \cap I_{j^*}$, $S \cup \{x\}$ should be shattered by \mathcal{T} according to our definition. Therefore $|I| \geq |I_{i^*} \cup I_{j^*}| + |I_{i^*} \cap I_{j^*}| = |I_{i^*}| + |I_{j^*}|$. In other words, $t(k, \epsilon, m, r) \geq 2t(k, \epsilon, m, \frac{\epsilon r}{\ell \binom{k}{2}})$. By induction the lemma is proved. \square

Proof of Lemma 1.4: Combining Lemma A.6 and Observation A.2, we get Lemma 1.4. \square

B Omitted Proofs from Section 4

Proof of Lemma 4.3: This is shown by a probabilistic argument. We randomly sample a number of covering k -partitions in the following way. Each time we sample a partition, we decide for each item in M , uniformly at random, which one of the k subsets it should be placed in. We repeat this process n times, and get a set of k -partitions $\{(T_1^i, T_2^i, \dots, T_k^i) \mid i \in [n]\}$. Let $A = \{a_1, a_2, \dots, a_\ell\} \subseteq [n]$, $B = \{b_1, b_2, \dots, b_\ell\} \in [k]^\ell$, let I_A^B denote the event

$$|T_{b_1}^{a_1} \cap T_{b_2}^{a_2} \cap \dots \cap T_{b_\ell}^{a_\ell}| > \frac{1}{k^\ell}(1 + \epsilon)m. \quad (6)$$

The expectation of the left hand side of (6) is m/k^ℓ . By Chernoff bound,

$$\Pr[I_A^B] \leq e^{-\epsilon^2 m/4k^\ell}, \quad \forall A, B.$$

The probability that I_A^B happens for some A and B is upper bounded by

$$\sum_{A \subseteq [n], |A|=\ell} \sum_{B \in [k]^\ell} \Pr[I_A^B] \leq \binom{n}{\ell} k^\ell e^{-\epsilon^2 m/4k^\ell} \leq \left(\frac{ken}{\ell}\right)^\ell e^{-\epsilon^2 m/4k^\ell}.$$

Therefore as long as $n < \frac{\ell}{ke} e^{\epsilon^2 m/4\ell k^\ell}$, the probability above is smaller than 1, i.e., there exist n partitions that satisfy the lemma. This completes the proof. \square

In the proof of Lemma 4.4, we used two forms of the Chernoff bound:

Theorem B.1. (*Chernoff bound*): Let X_1, X_2, \dots, X_n be i.i.d. random variables such that $X_i \in \{0, 1\}$ and $\Pr(X_i = 1) = p$ for every $i \in [n]$. Let $X = \sum_i X_i$ and $\mu = E[X]$, then

- (i) for $\delta > 2e - 1$, $\Pr[X > (1 + \delta)\mu] < 2^{-\delta\mu}$;
- (ii) for $0 < \delta < 2e - 1$, $\Pr[X > (1 + \delta)\mu] < e^{-\delta^2\mu/4}$.

C Omitted Proofs from Section 2.4

Proof of Lemma 2.7: Suppose S_1, \dots, S_k are CD sets, with circuit families $\{\mathcal{C}_n^{(i)}\}$ ($1 \leq i \leq k$) such that $\mathcal{C}_n^{(i)}$ has size bounded by a polynomial $q_i(n)$ and decides 3SAT correctly on all instances of size $n \in S_i$. Let $q(n)$ be a polynomial satisfying $q(n) \geq \max_{1 \leq i \leq k} q_i(n)$ for all $n \in \mathbb{N}$. We can obtain a family of circuits $\{\mathcal{C}_n\}$ of size bounded by $q(n)$, by defining \mathcal{C}_n to be equal to $\mathcal{C}_n^{(i)}$ if n belongs to S_i but not to S_1, \dots, S_{i-1} , and defining \mathcal{C}_n to be arbitrary if $n \notin S_1 \cup \dots \cup S_k$. Then \mathcal{C}_n decides 3SAT correctly on all instances of size $n \in S_1 \cup \dots \cup S_k$, as desired.

If S_1, \dots, S_k are CD sets, p_1, \dots, p_k are polynomials, and for $1 \leq i \leq k$ we have a PCD set $T_i \subseteq \bigcup_{n \in S_i} [n, p_i(n)]$, then we may take $p(n)$ to be any polynomial satisfying $p(n) \geq \max_{1 \leq i \leq k} p_i(n)$ for all $n \in \mathbb{N}$, and we may take S to be the set $S_1 \cup \dots \cup S_k$. Then we find that the set $T = T_1 \cup \dots \cup T_k$ is contained in $\bigcup_{n \in S} [n, p(n)]$. This implies that T is PCD, because S is CD. \square

Proof of Lemma 2.9: By our assumption that \mathcal{L} is NP-hard under polynomial-time many-one reductions, there is such a reduction from 3SAT to \mathcal{L} . Since the running time of the reduction is bounded by a polynomial $p(n)$, we know that it transforms a 3SAT instance of size n into an \mathcal{L} instance of size at most $p(n)$. Assume without loss of generality that $p(n)$ is an increasing function of n .

Let S be the set of all n such that $\{p(n) + 1, p(n) + 2, \dots, p(n + 1)\}$ intersects T . The set S is complexity-defying, because for any $n \in S$ we can construct a polynomial-sized circuit that correctly decides 3SAT instances of size n , as follows. First, we take the given 3SAT instance and apply the reduction from the preceding paragraph to transform it into an \mathcal{L} instance of size at most $p(n)$. Then, letting m be any element of $T \cap \{p(n) + 1, \dots, p(n + 1)\}$, we apply the padding reduction to transform this \mathcal{L} instance into another \mathcal{L} instance of size m . Finally, we solve this instance using a circuit of size $\text{poly}(m)$ that correctly decides \mathcal{L} on all instances of size m ; such a circuit exists by our assumption on T .

For every $m \in T$ there is an $n \in \mathbb{N}$ such that $p(n) < m \leq p(n + 1)$, and this n belongs to S . Thus, $T \subseteq \bigcup_{n \in S} [n, p(n + 1)]$, and this confirms that T is PCD. \square

Proof of Lemma 2.10: Suppose that

$$\mathbb{N} \subseteq \bigcup_{n \in S} [n, p(n)] \tag{7}$$

for some complexity-defying set S and polynomial function $p(n)$. We may assume without loss of generality that $p(n)$ is an increasing function of n and that $p(n) \geq n$ for all n .

Suppose that $\{\mathcal{C}_n\}$ is a polynomial-sized circuit family that correctly decides 3SAT whenever the input size is in S . We will construct a polynomial-sized circuit family that correctly decides 3SAT on all inputs. The construction is as follows: given an input size m , using (7) we may find a natural number n such that $n \leq p(m) \leq p(n)$. Since p is an increasing function, we know that $n \geq m$. Given an instance of 3SAT of size m , we first adjoin irrelevant clauses that don't affect its satisfiability — e.g. the clause $(x \vee \bar{x})$ — until the input size is increased to n . This transformation can be done by a circuit of size $\text{poly}(m)$, since $n \leq p(m)$. Then we solve the new 3SAT instance using the circuit \mathcal{C}_n . By our assumption on S , this correctly decides the original 3SAT instance of size m . As m was arbitrary, this establishes that $\text{NP} \subseteq \text{P/poly}$, as desired. \square

D Additional Preliminaries

D.1 Truthfulness

A k -bidder, m -item mechanism for combinatorial auctions with valuations in \mathcal{C} is a pair (f, p) where $f : \mathcal{C}_m^k \rightarrow \mathcal{X}([m], [k])$ is an *allocation rule*, and $p = (p_1, \dots, p_k)$ where $p_i : \mathcal{C}_m^k \rightarrow \mathbb{R}$ is a *payment scheme*. (f, p) might be either randomized or deterministic.

We say deterministic mechanism (f, p) is *truthful* if for all i , all v_i, v'_i and all v_{-i} we have that $v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v'_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i})$. A randomized mechanism (f, p) is *universally truthful* if it is a probability distribution over truthful deterministic mechanisms. More generally, (f, p) is *truthful in expectation* if for all i , all v_i, v'_i and all v_{-i} we have that $E[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \geq E[v'_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i})]$, where the expectation is taken over the internal random coins of the algorithm.

D.2 Algorithms and Approximation

Fix a valuation class \mathcal{C} . An algorithm \mathcal{A} for combinatorial auctions with \mathcal{C} valuations takes as input the number of players k , the number of items m , and a player valuation profile v_1, \dots, v_k where $v_i \in \mathcal{C}_m$. \mathcal{A} must then output an allocation of $[m]$ to $[k]$. For each k and m , \mathcal{A} induces an allocation rule of m items to k bidders. Our approximation bounds are all in terms of the number of players. Therefore, in our proofs we consider combinatorial auctions with a fixed number of bidders k .

We say an algorithm \mathcal{A} for k -player combinatorial auctions *achieves an α -approximation* if, for every input m and v_1, \dots, v_k :

$$E[v(\mathcal{A}(m, v_1, \dots, v_k))] \geq \alpha \max_{S \in \mathcal{X}([m], [k])} v(S)$$

Moreover, we say \mathcal{A} *achieves an α -approximation for m items* if the above holds whenever the number of items is fixed at m .

D.3 A Primer on Non-Uniform Computation

Non-uniform computation is a standard notion from complexity theory (see e.g. [3]). We say an algorithm is non-uniform if it takes in an extra parameter, often referred to as an *advice string*. However, the advice string is allowed to vary only with the size of the input (i.e. with m). Moreover, the length of the advice string can grow only polynomially in the size of the input. If a problem admits a non-uniform polynomial-time algorithm, this is equivalent to the existence of a family of polynomial-sized boolean circuits for the problem. When we say a non-uniform algorithm is polynomial-time MIWR, we mean that the algorithm runs in time polynomial in m , and maximizes over a weighted range, regardless of the advice string. When we say a non-uniform algorithm achieves an approximation ratio of α on m , we mean that there exists a choice of advice string for input length m such that the algorithm always outputs an α -approximate allocation.