Limitations of Randomized Mechanisms for Combinatorial Auctions

Shaddin Dughmi* Jan Vondrák†

January 9, 2014

Abstract

We address the following fundamental question in the area of incentive-compatible mechanism design: Are truthful-in-expectation mechanisms compatible with polynomial-time approximation? In particular, can polynomial-time truthful-in-expectation mechanisms achieve a near-optimal approximation ratio for combinatorial auctions with submodular valuations?

We prove that this is not the case: 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. In contrast, a non-truthful (1-1/e)-approximation algorithm is known [39], and a truthful-in-expectation (1-1/e)-approximation mechanism was recently developed for the special case of coverage valuations [20]. We also prove an analogous result for the combinatorial public projects (CPP) problem. Both our results present a significant separation between coverage functions and submodular functions, which does not occur for these problems without strategic considerations.

*University of Southern California, 941 Bloom Walk, Los Angeles, CA 90089. Email: shaddin@usc.edu. This work was done while the author was a student at Stanford University.

†IBM Almaden Research Center, 650 Harry Rd, San Jose, CA. E-mail: jvondrak@us.ibm.com.

1 Introduction

The overarching goal of *algorithmic mechanism design* is to design computationally-efficient mechanisms that solve or approximate fundamental resource allocation problems in which the underlying data is elicited from self-interested participants in the system. Work in this field has revealed a fundamental tension between the two main design goals in this domain, incentive-compatibility and computational efficiency. Understanding the power of mechanisms that satisfy both desiderata, in terms of their ability to approximate optimal allocations of resources, has therefore spawned a large literature of both positive and negative results that draw on ideas from algorithm design, computational complexity, and mechanism design.

In this paper, we consider mechanisms for *combinatorial auctions*. In combinatorial auctions, there is a set of items up for sale, and a set of self-interested players each of whom is equipped with a private *valuation function* mapping bundles of items to the player's value. The valuation functions are monotone non-decreasing, and normalized so that the value for the empty set is zero. We adopt the perspective of an auctioneer whose goal is to maximize social welfare, which in this context is the sum of the players' values for the bundles they receive. As is traditional in algorithmic mechanism design more generally, we examine the trade-off between incentive compatibility and computational efficiency[1] in this setting. Since we do not assume the existence of a prior on players' valuations, we focus on mechanisms that are dominant-strategy incentive compatible.

Combinatorial auctions enjoy paradigmatic status in mechanism design. We quote Blumrosen and Nisan [4]: "Combinatorial auctions serve as a common abstraction for many resource allocation problems in decentralized computerized systems such as the Internet, and may serve as a central building block of future electronic commerce systems." It is therefore unsurprising that combinatorial auctions have already been applied or considered in many contexts, such as in allocation of electromagnetic spectrum, allocation of airport take-off and landing slots, and more. (See Cramton et al [9].)

In many existing and potential applications of combinatorial auctions, the welfare-maximizing and incentive-compatible VCG mechanism cannot be deployed, due in part to the computational intractability of the welfare maximization problem. Combinatorial auctions that are actually employed, for instance in the practical settings mentioned above, are often heuristic in nature, specifically tailored to particular markets, and rigorous guarantees on their performance are rare in all but the simplest of settings. This has motivated a major research direction in algorithmic mechanism design, seeking a rigorous understanding of the space of computationally-efficient incentive-compatible mechanisms for variants of combinatorial auctions.

Strong impossibility results based on complexity-theoretic conjectures such as $P \neq NP$ [2] rule out such mechanisms with useful approximation guarantees in general.[3] Therefore, the community has focused on well-motivated special cases, by assuming some structure on the space of player valuations. Despite intense study over the past decade, however, constant-factor approximation mechanisms for combinatorial auctions have eluded researchers, even for restricted classes of valuations for which (non-incentive-compatible) constant-factor approximation algorithms are known. The most studied such variant, and until recently the most promising candidate for positive results, assumes that player valuations are *submodular*.[4]

There is extensive literature on welfare maximization in combinatorial auctions with submodular valuations. The state of the art in the non-strategic setting, i.e. when incentive-compatibility is dropped as

[1] We adopt the standard notion of efficiency used in computer science, namely *polynomial time*. For combinatorial auctions, we say an algorithm runs in polynomial time if it terminates after a number of steps that is polynomial in the number of players and items in the auction.

[2] The class P denotes computational problems that can be solved exactly in time polynomial in the representation size of their input. NP denotes problems for which a correct solution can be *verified* in polynomial time. The class P is known to be contained in NP, and the question of whether this containment is strict is the major open question of computational complexity.

[3] Like much of the related literature, we focus on worst-case multiplicative approximation guarantees in this work. An algorithm or mechanism for combinatorial auctions has an *approximation ratio* of α if it chooses an allocation whose welfare is at least a $1/\alpha$ fraction of the maximum possible for the reported valuations.

[4] Loosely speaking, submodular functions are those set functions that satisfy *diminishing marginal returns*. We define them formally in the preliminaries section.

a design requirement, is an $(\frac{e}{e-1})$-approximation algorithm due to Vondrák [39]. A matching impossibility result [27] shows this guarantee to be the best possible, assuming $P \neq NP$. Whereas the latter impossibility result applies to a specific class of submodular function with a certain representation, the algorithmic result of [39] holds under very general assumptions on the submodular functions and their representation — namely, that the algorithm has access to a *value oracle* which can be queried for a player's value for any particular bundle. Such an oracle model is a typical abstraction for identifying a class of computational problems of similar complexity, in this case the class of all combinatorial auction problems where players are equipped with valuation functions that are submodular and can be evaluated efficiently, regardless of their representation.[5]

When *both* incentive compatibility and computational efficiency are sought, the outlook for positive results has been more grim. A series of works have provided evidence that computational efficiency and incentive-compatibility are at loggerheads in combinatorial auctions. These works typically make assumptions about the sought mechanisms in three different dimensions: (1) The class of player valuations, most commonly submodular valuations or a close relative;[6] (2) the representation of the valuation functions, either via an oracle model or a chosen explicit representation; and (3) the notion of incentive-compatibility sought, with a focus on how it relates to the use of randomization in the mechanism's allocation and payment rules. For (3), three different restrictions of dominant-strategy incentive compatibility are commonly considered: (a) *truthfulness in expectation* allows a mechanism to randomize its choice of allocation and payments in the interest of computational efficiency, subject to the requirement that truthful reporting maximizes a risk-neutral player's expected payoff regardless of the reports of others;[7] (b) *universal truthfulness* makes no assumptions regarding players' risk attitudes,[8] and requires truth-telling to maximize a player's payoff for *every* draw of the mechanism's internal coins; and (c) *deterministic truthfulness* disallows the use of randomness in the mechanism entirely. Whereas (a),(b) and (c) are progressively more restrictive, the motivation for studying all three is born both of the traditional desire in computer science for understanding the power of randomization as a resource, as well as the fact that theorems regarding mechanisms that restrict the use of randomization have been more technically accessible.

After specifying each of these three parameters, a typical impossibility result rules out polynomial-time and incentive-compatible mechanisms achieving a certain worst-case approximation ratio, sometimes after making additional assumptions on the techniques used to design the mechanism. The recent result of Dobzinski [16] is the most relevant to our work, and proves that there is no universally-truthful mechanism for submodular combinatorial auctions in the value oracle model achieving an approximation ratio of $m^{\frac{1}{2}-\varepsilon}$, when m denotes the number of items in the auction and $\varepsilon > 0$ is independent of m. We overview other results along these lines in the related work section.

In light of the overwhelming body of negative results, it came as a surprise when a $(\frac{e}{e-1})$-approximate randomized mechanism was recently discovered by Dughmi, Roughgarden and Yan [20] for a large subclass of submodular valuations. Their mechanism applies to the canonical examples of submodularity, known as *coverage functions*,[9] as well as the larger family of *matroid rank sum functions*[10] in a particular oracle model they define. This development shed doubt on the importance of the existing body of negative results in the field: are these negative results merely a consequence of a restriction on the use of randomness? Can relaxing our solution concept to the traditional definition employed in much of economic theory, namely

[5]Other oracle models, such as the *demand oracle* model, or the *communication complexity* model, make stronger assumptions. We refer to the preliminaries section for definitions.

[6]Examples include sub-additive valuations, XOS valuations, and others. For a catalogue of valuations considered in combinatorial auctions, we refer the reader to [3].

[7]We note that truthfulness in expectation is the traditional notion of incentive compatibility considered in much of the economic literature.

[8]Indeed, universal truthfulness goes even further by not requiring that players are expected utility maximizers in the Von-Neumann Morgenstern sense.

[9]Loosely speaking, a *coverage function* defines for each item a subset of a measure space, and associates to each set of items a value equal to the measure of the union of the associated subsets. We define coverage functions formally in the preliminaries section.

[10]This is a family of set functions defined by a family of combinatorial objects known as *matroids*. We refer the reader to the preliminaries section for a formal definition.

truthfulness-in-expectation, be the cure for combinatorial auctions, perhaps enabling an optimal $(\frac{e}{e-1})$-approximation mechanism for all submodular valuations? There is a noteworthy historical precedent from the non-strategic variant of the same problem: the $(\frac{e}{e-1})$-approximation algorithm of Vondrak [39] was preceded by the analogous result for weighted matroid rank sum functions [7], which came after the same guarantee for coverage functions [14]. It is reasonable to anticipate a similar chain of events for mechanism design in the same setting.

Our results. We prove that this is not the case, and there is a significant separation between the classes of weighted matroid rank sums and general monotone submodular functions. More precisely, there is no truthful-in-expectation mechanism for combinatorial auctions with monotone submodular valuations in the value oracle model, guaranteeing an approximation better than $1/m^\gamma$ for some fixed $\gamma > 0$ (Theorem 5.1), where m denotes the number of items in the auction. In particular, the results of [20] cannot be extended to all monotone submodular valuations, provided that the valuations can be accessed only via value queries. (The possibility of an extension to submodular valuations given by an explicit representation was addressed in follow-up work which we discuss below.)

We also provide evidence that our impossibility result is robust in a certain sense. Specifically, we show that Thereom 5.1 holds even if we relax to an approximate notion of truthfulness in expectation. We are motivated by the fact that the result of [20] holds in a stronger oracle model than the value oracles we consider in our result. We show in Appendix B that the mechanism of [20] for weighted sums of matroid rank functions can be implemented in the value oracle model, at the cost of relaxing the solution concept to approximate truthfulness in expectation.[11] This stands in direct contrast to the case of monotone submodular valuations, where our result rules out such a mechanism.

We also prove a similar result for a problem that is related to combinatorial auctions both historically and technically: *flexible submodular combinatorial public projects*. We refer the reader to Section 2.2 for a formal definition of this problem, and to Section 4 for a brief history. We show that there is no polynomial-time, truthful-in-expectation mechanism providing an approximation better than $1/m^\gamma$ for some fixed $\gamma > 0$, when m denotes the number of public projects considered. This is true even in the case of a single player, and even after relaxing to approximate truthfulness. The combinatorial public projects problem admits simpler structure than combinatorial auctions, and hence we use it as a warm-up to demonstrate our approach.

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

Figure 1: Currently known results for combinatorial auctions: approximation | 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.

Related Work Mechanisms for combinatorial auctions have been studied by researchers in many disciplines, and the literature on them is vast both in its depth and in the breadth of the design constraints and environments considered. We can only hope to provide a glimpse of this body of work, biased as it may be to placing our work in the proper context. We begin with a brief outline of some practical applications of combinatorial auctions, and then proceed to overview existing work in the design of approximation mechanisms for variants of the problem. For a more comprehensive overview of research on combinatorial auctions,

[11] This result first appeared in [21], in joint work with Tim Roughgarden and Qiqi Yan. We include it here for completeness.

we refer the reader to references by Cramton et al. [9] and Milgrom [32], and a survey by Blumrosen and Nisan [3].

Combinatorial auctions have been considered and/or implemented in many contexts. The most prominent application is to allocation of the electromagnetic spectrum: many spectrum auctions have been employed by governments around the world since the mid 1990s, resulting in the sale of hundreds of billions of dollars worth of spectrum licenses (see [32]). Other applications include airport take-off and landing slot allocation, bus route allocation, and industrial procurement (see e.g. [9]).

The VCG mechanism is not employed, or even considered seriously, for many of the applications mentioned in the preceding discussion. This is in-part due to the computational intractability of its allocation rule, among other critiques.[12] Work in algorithmic mechanism design, like the results of this paper, isolates the first critique of the welfare-maximizing auction, and proposes the design of polynomial-time mechanisms with rigorous approximation guarantees as an alternative. This contrasts with much of the remaining literature and existing implementations of combinatorial auctions, which forgo rigorous guarantees on the quality of the allocation in favor of auctions with other desirable properties, such as simplicity, resistance to collusion, and empirical performance in specific markets. (See [4], [3], and [32] for a discussion.)

It is traditional to describe variants of combinatorial auctions via a class of valuations and an oracle model describing how these valuations are accessed. Many combinations of valuation class and oracle model have been considered, and we refer the interested reader to Figure 1 for an overview of some of the known results. A more comprehensive, though dated, overview appears in [3, Figure 11.2]. Until recently, there were no incentive-compatible constant-factor approximation mechanisms for combinatorial auctions with heterogeneous goods, even for settings where constant-factor approximation algorithms exist. The most promising and natural class of valuations for which such a positive result has been pursued is the class of submodular valuations, where a (non-incentive-compatible) $(\frac{e}{e-1})$-approximation algorithm exists [39]. The best incentive-compatible mechanisms which apply to all submodular valuations are a truthful-in-expectation $O\left(\frac{\log m}{\log\log m}\right)$-approximation mechanism in the communication complexity model due to Dobzinski et al. [15], and a universally-truthful $O(\log m \log\log m)$-approximation mechanism in the demand oracle model due to Dobzinski [10].

A series of works have provided evidence that computational efficiency and universal truthfulness are in conflict for combinatorial auctions. Universally-truthful constant-approximation mechanisms that simply adapt the VCG mechanism by restricting its range of allocations[13] have been ruled out for submodular combinatorial auctions in all oracle models[14] [12], and also for explicitly given sub-classes of submodular valuations assuming widely-believed complexity-theoretic conjectures [5]. In recent work, Dobzinski [16] proved that, in the value oracle model, there is no universally-truthful and polynomial-time mechanism for submodular combinatorial auctions achieving an approximation ratio better than $m^{\epsilon-1/2}$, where m denotes the number of items in the auction and $\varepsilon > 0$ is independent of m.

These results suggested that relaxation to truthfulness in expectation holds the only remaining promise for positive results. A recent result by Dughmi, Roughgarden and Yan [20] buoyed this hope by designing a truthful-in-expectation $(\frac{e}{e-1})$-approximation mechanism for a large subclass of submodular valuations. Their mechanism applies to explicitly represented *coverage functions* – the canonical examples of submodular functions. More generally, their mechanism applies to "black-box" valuations that are expressible as weighted sums of *matroid rank functions*, provided they support "lottery-value queries" (what is the expected value $\mathbf{E}[v_i(\widehat{\mathbf{x}})]$ for a given product distribution $\widehat{\mathbf{x}}$). Our impossibility result can be viewed as an extension of that of [16], ruling out generalization of the mechanism of [20] to all submodular functions.

[12] The VCG mechanism has other weaknesses that have been criticized in a combinatorial auctions context, including vulnerability to collusion, low seller revenues in some settings, and more. See Ausubel and Milgrom [2] for a discussion.

[13] These mechanisms have been referred to as *maximal-in-range*, and have dominated the design of deterministic polynomial-time approximation mechanisms in much of the recent literature. We define them formally, along with their randomized analogues, in the preliminaries section.

[14] Specifically, they have been ruled out in the *communication complexity model*, defined formally in the preliminaries section.

Follow-up work. While our impossibility result rules out truthful-in-expectation mechanisms in the value oracle model, subsequent work has addressed the possibility of designing mechanisms for valuations supplied in some compact explicit form. More precisely, the question here is whether it is possible to design a mechanism for combinatorial auctions when each agent presents an explicit description of her valuation v_i in some agreed-upon format, using space polynomial in the number of items m. It turns out that this does not help, in the sense that the impossibility results in this paper essentially carry over to the case of monotone submodular valuations represented by a certain compact encoding. Techniques for translating impossibility results in the value oracle model to the compact representation model were developed in [17] and then applied to combinatorial auctions in [18]. This work proves that even for (a certain subclass of) monotone sumodular valuations which can be represented compactly, there is no truthful-in-expectation mechanism achieving approximation better than m^{γ} for some constant $\gamma > 0$, unless $NP \subseteq P/poly$.

Our techniques. Our hardness results are obtained by combining two recently developed techniques: the *symmetry gap* technique for submodular functions [40], and the *direct hardness* approach for combinatorial auctions [16].

First, we consider the possibility of maximal-in-distributional range (MIDR) mechanisms.[15] We endeavor to explain why the approach of [20] 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 [20]. 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*[16] 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 — in terms of approximating combinatorial auctions using a polynomial number of value queries — 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 [16] 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 to obtain a hardness of $m^{-\gamma}$-approximation.

Organization of the paper. After the necessary preliminaries (Section 2), we present our intuition on the separation between coverage valuations and submodular valuations in Section 3. In Section 4, we present the proof of hardness for combinatorial public projects, and in Section 5 the proof for combinatorial auctions. In Appendix A, we present a transformation from approximately-TIE to approximately-MIDR mechanisms for combinatorial public projects. In Appendix B, we present a transformation from approximately-MIDR to approximately-TIE mechanisms for combinatorial auctions; in particular, an approximately-TIE variant of the mechanism of [20] in the value oracle model. Additional technical lemmas are deferred to appendices as well.

Acknowledgment. We would like to thank Tim Roughgarden and Qiqi Yan for the suggestion that we include the results of Appendix B, which appeared previously in [21], in this paper.

[15] These mechanisms are adaptations of VCG that restrict the space of allocations in the interest of reducting computational complexity, and are the primary tool used in many of the recent positive results in algorithmic mechanism design, including the result of [20]. We define them formally in the preliminaries section.

[16] An algorithm is non-uniform if it is supplemented with an *advice string* for each input size. Allowing such an advice string expands the set of efficiently computable functions.

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* $v_i : \Omega \to \mathbb{R}_+$. 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 $v = (v_1, \ldots, v_n)$ 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 $v = (v_1, \ldots, v_n)$, 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

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

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. [35].

Let $dist(\Omega)$ denote the probability distributions over the set of feasible solutions Ω, and let $\mathcal{R} \subseteq dist(\Omega)$ be a compact subset of them. The corresponding *Maximal in Distributional Range (MIDR)* allocation rule is defined as follows: given reported valuation functions $v_1, \ldots, v_n$, 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 [11]). 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

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

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.)

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

Approximately truthful mechanisms are related to *approximately maximal-in-distributional-range* allocation rules. An allocation rule $\mathcal{A} : \mathcal{V} \to \Omega$ is $(1-\epsilon)$-approximately maximal-in-distributional range if it fixes a

$\mathcal{R} \subseteq dist(\Omega)$, 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 [20, 22], if implemented in the value oracle model, can be made $(1-o(1))$-approximately truthful-in-expectation (see Appendix B). It is not known whether these mechanisms can be made truthful-in-expectation, and indeed the value oracle model might be too weak to achieve this. Our results provide a separation between the classes of weighted matroid rank sums and submodular functions, regardless of this issue.

2.2 Combinatorial Auctions and Public Projects

In *Combinatorial Auctions* there is a set M of m items, and a set of n players. Each player i has a valuation function $v_i : 2^M \to \mathbb{R}_+$ 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* $(S_1, \ldots, S_n)$, 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 $(S_1, \ldots, S_n)$ is equal to $v_i(S_i)$. The goal is to choose an allocation maximizing *social welfare*: $\sum_i v_i(S_i)$.

In *Combinatorial Public Projects* there is a set $[m] = \{1, \ldots, m\}$ of *projects*, a cardinality bound k such that $0 \leq k \leq m$, and a set $[n] = \{1, \ldots, n\}$ of *players*. Each player i has a valuation function $v_i : 2^{[m]} \to \mathbb{R}_+$ 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 [38, 6, 16] 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*: $\sum_i v_i(S)$.

2.3 Describing Variants of Combinatorial Auctions and Public Projects

Variants of combinatorial auctions and public projects are typically specified by two parameters: A *class of valuations* assumed to include each player's valuation function, and a corresponding *oracle model* that describes how the algorithm may access the valuation. As an example, the least tractable class of valuations typically considered, often referred to as the set of *unrestricted valuations*,[17] is the family of all functions from bundles to real numbers that are non-decreasing and normalized. Combinatorial auctions are most interesting to study when additional structure is assumed of the valuation functions, in which case we say the combinatorial auctions problem has *restricted valuations*.

Arguably the most studied class of restricted valuations is that of *submodular valuations*. A function $v : 2^{[m]}$ is submodular if it satisfies *diminishing marginal returns*: specifically, the marginal value $v(S \cup \{j\}) - v_i(S)$ for a each fixed item j is non-increasing in S. Other related classes include *subadditive valuations* (see e.g. [24]), XOS valuations (see e.g. [13]), among others.

Two sub-classes of submodular valuations that are particularly relevant to our work are *coverage valuations* and *matroid rank sum valuations*, as they are the main classes of valuations considered in [20]. A *coverage function* f on ground set $[m]$ designates some measure space L, and m subsets $A_1, \ldots, A_m \subseteq L$, such that $f(S)$ is the measure of $|\cup_{j \in S} A_j|$. A set function $v : 2^{[m]} \to \mathbb{R}_+$ is a *matroid rank sum (MRS)* function if there exists a family of matroid rank functions[18] $u_1, \ldots, u_\kappa : 2^{[m]} \to \mathbb{N}$, and associated non-negative weights $w_1, \ldots, w_\kappa \in \mathbb{R}^+$, such that $v(S) = \sum_{\ell=1}^{\kappa} w_\ell u_\ell(S)$ for all $S \subseteq [m]$. We note that all coverage functions are matroid rank sum functions, and both are submodular functions.

Another important class of valuation functions is that of *budget-additive valuations*: These are functions in the form $f(S) = \min\{\sum_{i \in S} a_i, B\}$. This is a subclass of submodular functions which is incomparable to

[17] Note that this term is overloaded. In context of a general mechanism design problem, an unrestricted valuation may depend arbitrarily on the allocation. On the other hand, an unrestricted valuation for combinatorial auctions must depend only on the player's own bundle, and be normalized and non-decreasing in said bundle, but may otherwise be arbitrary.

[18] Matroids, and their rank functions, are combinatorial objects that commonly appear in algorithm and mechanism design. Their formal definition is beyond the scope of this paper, and we refer the reader to [37].

the classes of coverage and matroid rank functions.

In general, a valuation function on m items is an exponential-sized object, requiring 2^m real numbers for explicit representation. This motivates various *oracle models* as alternatives to explicit representation. The most traditional such model is that of *value oracles*, where an algorithm may query a valuation function at any specific bundle in constant time. A more powerful model is that of *demand oracles*, where a valuation v is given by an oracle that takes as input a price p_j for each item j, and returns the bundle S maximizing the player's utility $v(S) - \sum_{j \in S} p_j$.[19] The most powerful access model of all is the *communication complexity model*, which assumes that an oracle for each valuation v answers arbitrary questions that can be formulated and answered using a polynomial number of bits. Whereas most natural examples of submodular functions allow for answering value oracle queries efficiently, the same can not be said of other, more powerful, oracle models.

It is also not uncommon to directly consider combinatorial auctions with valuations that are represented explicitly, rather than via an oracle. Valuation classes considered in this manner are typically defined in reference to a particular short representation; examples include single-minded valuations, where a player has non-zero value for a single bundle represented explicitly, and the class of coverage valuations when L is a finite set listed explicitly. When valuations are explicitly represented, the mechanism's runtime is allowed to depend polynomially on the size of the representation, as is traditional in algorithm design. This obviates the need for an oracle model, and allows traditional complexity-theoretic analysis of the resulting combinatorial allocation problem. On the other hand, the main advantage of an oracle model is that it enables positive results that make minimal assumptions on how valuations are represented, and identifies classes of problems of similar computational complexity. For example, a polynomial-time algorithm or mechanism for submodular combinatorial auctions in the value oracle model applies to any explicitly represented subclass of submodular valuations for which value queries can be computed efficiently.

2.4 A truthful-in-expectation mechanism for coverage valuations

Here we describe briefly a truthful-in-expectation mechanism, based on the VCG payment scheme (developed in [20]). This mechanism is particularly important in this paper. We assume here that each player submits a *coverage* valuation function, $f_i(S) = |\bigcup_{j \in S} A_{ij}|$, represented explicitly by subsets A_{ij} of some universe U.

The mechanism works as follows: Suppose the number of items is m. The valuation function for each player is formally replaced by a continuous function $F_i^{exp}(x_{i1}, \ldots, x_{im}) = \sum_{a \in U} \left(1 - e^{-\sum_{j: a \in A_{ij}} x_{ij}}\right)$. This function has a natural interpretation: It gives the expected number of elements of U covered if player i receives each item j independently with probability $1 - e^{-x_{ij}}$. The mechanism then finds a *fractional allocation* x_{ij} (satisfying $x_{ij} \in [0,1]$ and $\sum_{j=1}^{m} x_{ij} = 1$ for each i), that maximizes the objective function $\sum_{i=1}^{n} F_i^{exp}(x_{i1}, \ldots, x_{im})$. Since this is a concave maximization problem, it can be solved efficiently.

Given this fractional allocation, items are allocated randomly as follows: player i receives item j with probability $1 - e^{-x_{ij}}$, independently between items. This is possible, since $\sum_{i=1}^{n}(1 - e^{-x_{ij}}) \leq \sum_{i=1}^{n} x_{ij} = 1$. Finally, the mechanism charges payments defined by the VCG scheme: each player pays the damage that he incurs to the remaining players in the above optimization problem. This can be shown to be a truthful-in-expectation mechanism (see [20] for more details).

3 Intuition - what fails for submodular valuations

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 [20]. 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.

[19] Demand queries are NP-hard to answer even for simple submodular functions such as budget-additive functions [31].

Let us recall the main idea of [20]: Let $f : 2^M \to \mathbb{R}_+$ 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 [7, 39, 28, 40, 36]. 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 [20] 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 rounding scheme to a point $\mathbf{x}$ is measured by a modified function $F^{exp}(x_1, \ldots, x_m) = F(1 - e^{-x_1}, \ldots, 1 - e^{-x_m})$, 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 [20] 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 [25, 33, 40]. The heart of a symmetry gap argument is a construction of a submodular function which is symmetric under certain transformations (such as permutations of coordinates), and it has the property that *symmetric solutions* are substantially worse than *asymmetric solutions*. The particular type of function we are interested in here is the function F^{exp} derived from a discrete submodular function. Our goal is to construct a function F^{exp} such that there is a significant gap between its value on symmetric vs. asymmetric fractional solutions.

The budget-additive function above is not very suitable for a symmetry gap construction, 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 $f_1, f_2 : 2^M \to [0,1]$ are monotone submodular functions, then

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

is also a monotone submodular function (see Lemma E.1). 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. Then we set

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

Here, $(y)_+ = \max\{y, 0\}$ denotes the positive part of a number. By Lemma E.1, $f(S)$ is a monotone submodular function. Let's consider the function $F^{exp}(x_1, \ldots, x_{2m}) = F(1 - e^{-x_1}, \ldots, 1 - e^{-x_{2m}})$. If $\sum_{i=1}^m x_i = \Omega(m)$ and $m \to \infty$, a random set obtained by sampling with probabilities $1 - e^{-x_i}$ will have cardinality very close to $\sum(1 - e^{-x_i})$ (by a Chernoff bound). We obtain

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

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 $F^{exp}\left(\frac{1}{2}\mathbf{1}_M\right) \simeq 1 - (-1 + 2e^{-1/2})^2 = 4e^{-1/2} - 4e^{-1} \simeq 0.955$. 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 [40], building up on previous work [25, 33], 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 [25, 33, 40], 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 [38] 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 [38]. Using this characterization, it was proved in [38] 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 succintly 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 [34]. 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 [6]: 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 [16] 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 [16], 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, [16] 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 [30, 11, 19, 20]. Flexible CPP has exhibited a similar pattern; Dughmi [22] 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 [23]), 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 [38, 6], 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 the 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 depends on m and δ, but given $\delta > 0$ it can be fixed for each input size m and acts as an "advice string" to the mechanism. We present the proof in Appendix A.

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, using poly(m) value queries (for any polynomial in m), and 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 poly(m).*

In the following, we present a sketch of the proof of this theorem.

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, corresponding to different possible inputs to the mechanism. (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 $D_0 \in \mathcal{R}$ such that a set sampled from it contains at least a c-fraction of $A^{(0)}$ in expectation. 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).

At level j, $1 \leq j \leq \ell$, we consider 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 provide better expected value than the distribution D_{j-1} guaranteed to be in $\mathcal{R}$ by induction from the previous level. (By density, we mean a certain notion of average size for sets sampled from D_j.) However, D_{j-1} depends on the set $A^{(j-1)}$ which was revealed to the mechanism on the previous level. 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 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.

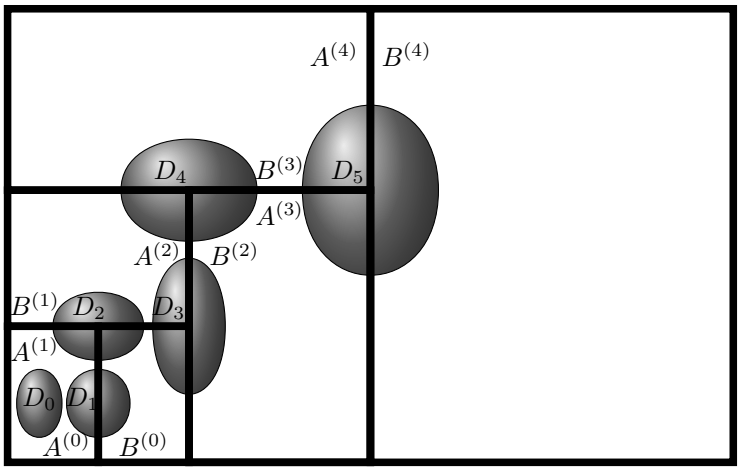

Figure 2: A bisection sequence $(A^{(j)}, B^{(j)})$, with the distributions D_j returned by the mechanism at level j. The density of D_j increases in a certain technical sense exponentially in j, although much slower than 2^j.

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.

4.1 The basic setup

We consider a ground set of $|M| = m = 400^\ell$ items for some $\ell \geq 1$. We set the cardinality bound to be $k = 200^\ell = m/n$ where $n = 2^\ell$. (We note that n is just a parameter unrelated to the number of players, which is 1 in this case. This parameter will however denote the number of players in Section 5.) An important object in the following will be a *random bisection sequence*.

Definition 4.3. *A random bisection sequence is a random sequence of pairs of sets $(A^{(0)}, B^{(0)})$, $(A^{(1)}, B^{(1)})$, ..., $(A^{(\ell)}, B^{(\ell)})$ generated as follows. We define $A^{(\ell)} = B^{(\ell)} = M$. Given $A^{(j)}$ for $0 < j \leq \ell$, we pick $(A^{(j-1)}, B^{(j-1)})$ uniformly among all partitions of $A^{(j)}$ into two parts of size $\frac{1}{2}|A^{(j)}|$.*

I.e., $|A^{(j)}| = |B^{(j)}| = 2^{j-\ell}m$. We refer to $A^{(j)} = A^{(j-1)} \cup B^{(j-1)}$ as the j-th level of the bisection sequence. Observe that the distribution of $(A^{(j-1)}, B^{(j-1)})$ is uniform among all pairs of disjoint sets of size $2^{j-1-\ell}m$. We will use valuation functions associated with each level of a bisection sequence. We denote these valuation functions at level j by $f_{A^{(j-1)},B^{(j-1)}}$. In particular, this valuation function depends only on the elements of $A^{(j)} = A^{(j-1)} \cup B^{(j-1)}$.

The bisection sequence is generated at random and unknown to the mechanism. For each particular choice of a valuation function at a certain level, the mechanism needs to produce a probability distribution over feasible sets, which is purportedly the (approximately) optimal one over a certain fixed range of distributions $\mathcal{R}$. The distribution will depend on the choice of a valuation function, in particular on the relevant set of items $A^{(j)}$. A function assigning a distribution over sets to every set $A^{(j)}$ is a complicated object; in order to be able to argue about all possible such functions, we distill the important information into a single random variable for each level j.

Definition 4.4. *We say that a random variable X_j is constructible by a range $\mathcal{R}$ at level j, if there is a distribution $D(A^{(j)}) \in \mathcal{R}$ for each set $A^{(j)}$ of size $2^{j-\ell}m$ such that if a random set R is generated by first*

choosing $A^{(j)}$ uniformly among all sets of size $2^{j-\ell}m$ and then sampling R from the distribution $D(A^{(j)})$, then

$$X_j = \frac{|R \cap A^{(j)}|}{|A^{(j)}|}.$$

Note that the normalization is chosen so that we have $X_j \in [0,1]$. There are two sources of randomness in defining X_j: one is the randomness in $A^{(j)}$, and one arises from the probability distribution $D(A^{(j)})$.

The first useful fact is the following (easy) lemma.

Lemma 4.5. *Consider a mechanism returning distributions from a range $\mathcal{R}$ that achieves a c-approximation for the problem $\max\{f(S) : |S| \leq k\}$ for f monotone submodular. Then there is a random variable X_0 constructible by $\mathcal{R}$ at level 0 such that*

$$\mathbf{E}[X_0] \geq c.$$

Proof. Consider the valuation function $f_{A^{(0)}}(S) = \frac{|S \cap A^{(0)}|}{|A^{(0)}|}$. Let $X_0 = \frac{|R^{(0)} \cap A^{(0)}|}{|A^{(0)}|}$ where $R^{(0)}$ is the random set returned by the mechanism, given valuation $f_{A^{(0)}}$ for $A^{(0)}$ chosen randomly among all sets of size $2^{-\ell}m$. By Definition 4.4, X_0 is a random variable constructible by the range $\mathcal{R}$ at level 0.

Since the optimum under valuation $f_{A^{(0)}}$ is 1 (achieved by $A^{(0)}$ itself), the mechanism should return expected value at least c. The value returned by the mechanism is exactly the random variable X_0, hence $\mathbf{E}[X_0] \geq c$. □

We remark that Lemma 4.5 is the only place where we use the assumption of c-approximation. In the following, our goal is to use the MIDR property to argue about distributions that must be in the range at higher levels, and prove successive bounds on the random variables $X_1, X_2, \ldots$.

4.2 The symmetry gap argument

The main building block of our proof is a symmetry gap argument whose goal is to show the following. If the mechanism optimizes over a certain fixed range $\mathcal{R}$ which supports distributions of "high density" at level j, then $\mathcal{R}$ must support distributions of even higher density (when properly scaled) at level $j+1$. However, the way we measure density is quite intricate. Recall the random variables $X_0, X_1, \ldots, X_\ell$ that encode certain distributions in the range at each level. It would be nice to say that for any X_j constructible at level j, there must be X_{j+1} constructible at level $j+1$ such that $\mathbf{E}[X_{j+1}] > \frac{1+\delta}{2}\mathbf{E}[X_j]$ (which would correspond to sets of larger cardinality at level $j+1$ than j). But this is not true - the distributions of X_j, X_{j+1} also matter and we cannot get a guaranteed boost just in terms of expectation. Instead, we define our measure of density as $\mathbf{E}[\phi(X_j)]$, using a concave function ϕ that we specify later. The symmetry gap argument allows us to prove the following.

Lemma 4.6. *Let $\phi : [0,1] \to [0,1]$ be a non-decreasing concave function. Fix $j \in \{0, 1, \ldots, \ell-1\}$ and a set $A^{(j+1)}$ of size $2^{j+1-\ell}m \geq m/n$. Let $(A^{(j)}, B^{(j)})$ be a random partition of $A^{(j+1)}$ into two sets of equal size. Then there is a monotone submodular function $\widetilde{f}_{A^{(j)},B^{(j)}}$ for each partition $(A^{(j)}, B^{(j)})$ such that*

- *For any distribution of a random set $R^{(j)}$ (possibly correlated with $A^{(j)}$) and the associated random variable $X_j = \frac{|R^{(j)} \cap A^{(j)}|}{|A^{(j)}|}$, we have*

$$\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j)})] \geq \mathbf{E}[\phi((X_j - nm^{-1/2})_+)].$$

- *Any mechanism that uses poly(n) value queries, when applied to the random input $\widetilde{f}_{A^{(j)},B^{(j)}}$ will return a random set $R^{(j+1)}$ such that for the random variable $X_{j+1} = \frac{|R^{(j+1)} \cap A^{(j+1)}|}{|A^{(j+1)}|}$,*

$$\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j+1)})] \leq \mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] + e^{-\Omega(n)}.$$

The expectations are over both $(A^{(j)}, B^{(j)})$ and $R^{(j)}$ or $R^{(j+1)}$ respectively.

The key point here is that the performance of a mechanism depends only on the fraction of elements taken from $A^{(j+1)}$, and not on the partition $(A^{(j)}, B^{(j)})$. While there might be a "good distribution" $R^{(j)}$ in the range which is correlated with $A^{(j)}$, the mechanism cannot find such a distribution and must compensate for it by returning larger sets. This will be important later.

The proof of Lemma 4.6 relies on the notion of symmetry gap developed in [25, 33, 40]. Since what we need here is a special case where the construction can be carried out explicitly quite easily, we present a self-contained proof here instead of referring to the general framework of [40].

Proof. Consider a pair of sets $(A^{(j)}, B^{(j)})$. Given a non-decreasing concave function $\phi : [0,1] \to [0,1]$, we define the valuation function $f_{A^{(j)},B^{(j)}}$ as follows:

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

The function depends only on how many elements we take from $A^{(j)}$ and how many from $B^{(j)}$. Moreover, the two sets play the same role in $f_{A^{(j)},B^{(j)}}$; i.e., all elements in $A^{(j+1)} = A^{(j)} \cup B^{(j)}$ contribute equivalently to $f_{A^{(j)},B^{(j)}}$. This is the kind of situation where we can apply a symmetry gap argument.

Let us simplify the notation and write

$$\psi(x,y) = 1 - (1 - \phi(x))(1 - \phi(y)),$$

where $x, y \in [0,1]$; i.e. $f_{A^{(j)},B^{(j)}}(S) = \psi\left(\frac{|S\cap A^{(j)}|}{|A^{(j)}|}, \frac{|S\cap B^{(j)}|}{|B^{(j)}|}\right)$. It is elementary to verify that since ϕ is non-decreasing concave, the first partial derivatives of ψ are non-negative and non-increasing with respect to both coordinates. Now we replace ψ by a modified function $\widetilde{\psi}$ which has the property that if $|x-y|$ is very small, the function value depends only on $x+y$. This can be accomplished explicitly as follows: For some $\beta > 0$, let

- $\widetilde{\psi}(x,y) = \psi(\frac{1}{2}(x+y), \frac{1}{2}(x+y))$ if $|x-y| \le \beta$.
- $\widetilde{\psi}(x,y) = \psi(x - \frac{1}{2}\beta, y + \frac{1}{2}\beta)$ if $x - y > \beta$.
- $\widetilde{\psi}(x,y) = \psi(x + \frac{1}{2}\beta, y - \frac{1}{2}\beta)$ if $y - x > \beta$.

Geometrically, this construction can be seen as taking the graph of $\psi(x,y)$, pulling it away from the diagonal $x = y$ on both sides, and patching the area close to the diagonal with a function which depends only on $x + y$ and is equal to the function on the diagonal. Using the properties of ψ, one can check that again the first partial derivatives of $\widetilde{\psi}$ are non-negative and non-increasing with respect to both coordinates. We define the function promised by the lemma as

$$\widetilde{f}_{A^{(j)},B^{(j)}}(S) = \widetilde{\psi}\left(\frac{|S \cap A^{(j)}|}{|A^{(j)}|}, \frac{|S \cap B^{(j)}|}{|B^{(j)}|}\right).$$

The properties of $\widetilde{\psi}$ imply that $\widetilde{f}_{A^{(j)},B^{(j)}}$ is a monotone submodular function (see e.g. [33, 40]).

We observe the following (which is the case in all proofs using the symmetry gap). For a "typical query" S, oblivious to the random partition $(A^{(j)}, B^{(j)})$, with high probability S will contain approximately the same number of elements from these two sets. (Recall that $|A^{(j)}| = |B^{(j)}|$.) We call a query S balanced if the parameters $x = \frac{|S\cap A^{(j)}|}{|A^{(j)}|}$ and $y = \frac{|S\cap B^{(j)}|}{|B^{(j)}|}$ are in the range where $|x-y| \le \beta$, and hence $\widetilde{f}_{A^{(j)},B^{(j)}}(S) = \widetilde{\psi}(x,y) = \psi(\frac{1}{2}(x+y), \frac{1}{2}(x+y))$ is independent of the particular partition $(A^{(j)}, B^{(j)})$. By Lemma D.1 (applied to the ground set $A^{(j+1)}$), the probability that any fixed query S is unbalanced is exponentially small:

$$\Pr[|x-y| > \beta] = \Pr\left[||S \cap A^{(j)}| - |S \cap B^{(j)}|| > \beta |A^{(j)}|\right] \le e^{-\Omega(\beta^2 |A^{(j+1)}|)}.$$

Recall that $|A^{(j+1)}| = 2^{j+1-\ell} m \ge m/n$. Therefore, if we pick $\beta = nm^{-1/2}$, the probability is $e^{-\Omega(n)}$. Let us fix for now the random coin flips of the mechanism. As long as all query answers are independent of the

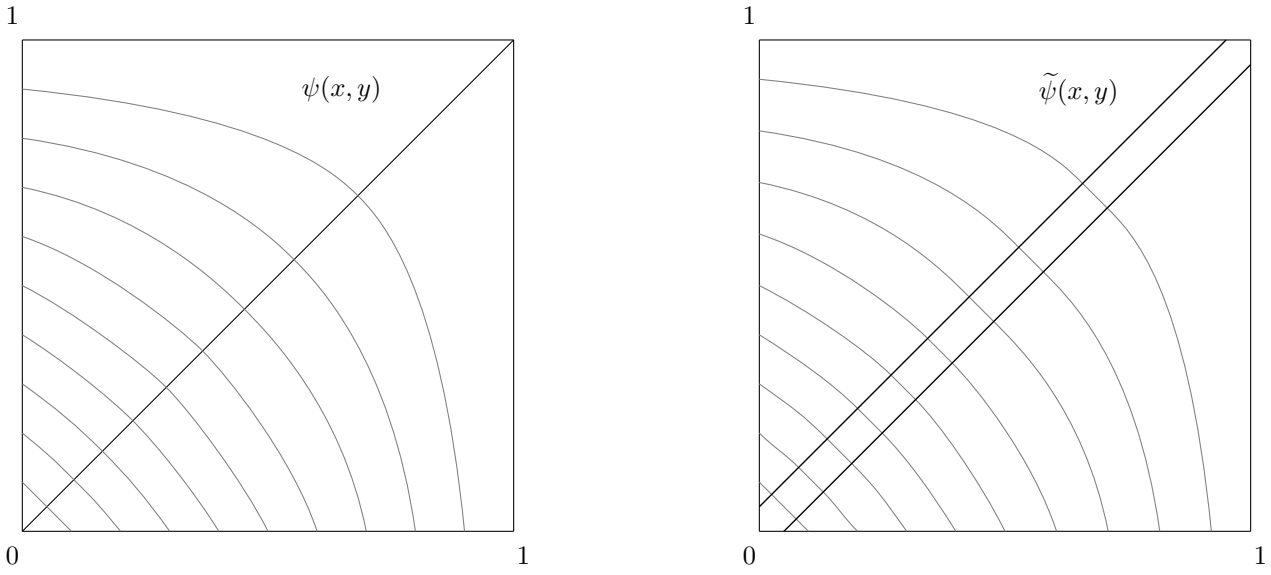

Figure 3: Construction of $\widetilde{\psi}(x,y)$ from $\psi(x,y)$, assuming that $\psi(x,y) = 1-(1-x)(1-y)$. The solid lines denote the diagonal $x=y$ and the shifted diagonals $x-y=\pm\beta$. The gray lines are the level sets of $\psi(x,y)$ and $\widetilde{\psi}(x,y)$.

partition $(A^{(j)}, B^{(j)})$, the mechanism will follow the same computation path, independent of $(A^{(j)}, B^{(j)})$, and we can use a union bound over its $poly(n)$ queries. Hence, the probability that a mechanism ever makes a query such that $|x-y| > \beta$ is $poly(n)e^{-\Omega(n)} = e^{-\Omega(n)}$. This is still true if we average over the random coin flips of the algorithm. Therefore, the output of the mechanism will be independent of $(A^{(j)}, B^{(j)})$ with probability $1 - e^{-\Omega(n)}$.

To summarize, the output of the mechanism, $R^{(j+1)}$, is with high probability independent of $(A^{(j)}, B^{(j)})$ and again by Lemma D.1 with high probability balanced with respect to $(A^{(j)}, B^{(j)})$. Given the definition of the random variable $X_{j+1} = \frac{|R^{(j+1)} \cap A^{(j+1)}|}{|A^{(j+1)}|}$, this means the output random set contains an X_{j+1}-fraction of the set $A^{(j+1)}$, approximately balanced between its two halves. For some $|\beta'| \le \frac{1}{2}\beta$, the value of such a set is

$$\widetilde{\psi}(X_{j+1} + \beta', X_{j+1} - \beta') = \psi(X_{j+1}, X_{j+1}) = 1 - (1 - \phi(X_{j+1}))^2.$$

Thus the expected value of this solution is $\mathbf{E}[f_{A^{(j)},B^{(j)}}(R^{(j+1)})] \le \mathbf{E}[1-(1-\phi(X_{j+1}))^2] + e^{-\Omega(n)}$ (where $e^{-\Omega(n)}$ accounts for the small probability of finding an unbalanced solution, whose value could be up to 1). This proves the second statement of the lemma.

Finally, consider any random set $R^{(j)}$ and the associated random variable $X_j = \frac{|R^{(j)} \cap A^{(j)}|}{|A^{(j)}|}$. We have

$$\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j)})) \ge \widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j)} \cap A^{(j)}) = \widetilde{\psi}(X_j, 0) \ge \psi\left((X_j - \beta)_+, 0\right) = \phi\left((X_j - \beta)_+\right).$$

Therefore $\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j)})] \ge \mathbf{E}[\phi((X_j - \beta)_+)]$. Recall that $\beta = nm^{-1/2}$, so this proves the first statement of the lemma. □

Considering the setup of random variables $X_0, X_1, \ldots, X_\ell$ constructible by $\mathcal{R}$ at different levels (Section 4.1), we obtain the following.

Lemma 4.7. *Consider a mechanism of polynomial query-complexity that $(1-\epsilon)$-approximately maximizes over a range of distributions $\mathcal{R}$ for the problem $\max\{f(S) : |S| \le k\}$ for f monotone submodular, ground set of size $m = 400^\ell$ and $k = 2^{-\ell}m$. Let $\phi : [0,1] \to [0,1]$ be a non-decreasing concave function. If a random*

variable X_j is constructible by $\mathcal{R}$ at level j, then there is a random variable X_{j+1} constructible by $\mathcal{R}$ at level $j+1$ such that

$$\mathbf{E}[1-(1-\phi(X_{j+1}))^2] \geq (1-\epsilon)\mathbf{E}[\phi((X_j - 10^{-\ell})_+)] - 10^{-\ell}.$$

Proof. Given ϕ, let $\widetilde{f}_{A^{(j)},B^{(j)}}$ be the valuation function provided by Lemma 4.6. Consider $A^{(j+1)}$ uniformly random among sets of size $2^{j+1-\ell}m$, bisected randomly into $A^{(j)} \cup B^{(j)}$. If X_j is constructible by the range $\mathcal{R}$ at level j, it means that for each $A^{(j)}$ there is a distribution $D(A^{(j)})$ in $\mathcal{R}$ such that $X_j = \frac{|R^{(j)} \cap A^{(j)}|}{|A^{(j)}|}$ where $A^{(j)}$ is random and $X^{(j)}$ is sampled from $D(A^{(j)})$. By Lemma 4.6, conditioned on any $A^{(j+1)}$ and taking expectation over the random partition $(A^{(j)}, B^{(j)})$, $\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j)}) \mid A^{(j+1)}] \geq \mathbf{E}[\phi((X_j - nm^{-1/2})_+) \mid A^{(j+1)}]$. Therefore the same holds also without the conditioning. Recall that we have $nm^{-1/2} = 2^{\ell}400^{-\ell/2} = 10^{-\ell}$. So we get

$$\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j)})] \geq \mathbf{E}[\phi((X_j - nm^{-1/2})_+)] = \mathbf{E}[\phi((X_j - 10^{-\ell})_+)].$$

Now let us run the mechanism on the same random instance and denote the output random set by $R^{(j+1)}$. By Lemma 4.6, $\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j+1)}) \mid A^{(j+1)}] \leq \mathbf{E}[1-(1-\phi(X_{j+1}))^2 \mid A^{(j+1)}] + e^{-\Omega(n)}$, where $X_{j+1} = \frac{|R^{(j+1)} \cap A^{(j+1)}|}{|A^{(j)}|}$. Hence this holds also without the conditioning:

$$\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j+1)})] \leq \mathbf{E}[1-(1-\phi(X_{j+1}))^2] + e^{-\Omega(n)} \leq \mathbf{E}[1-(1-\phi(X_{j+1}))^2] + 10^{-\ell}$$

and by definition X_{j+1} is constructible by $\mathcal{R}$ at level $j+1$.

To conclude, if the mechanism maximizes $(1-\epsilon)$-approximately over $\mathcal{R}$, then the expected value of $R^{(j+1)}$ conditioned on $(A^{(j)}, B^{(j)})$ must be at least $(1-\epsilon)\times$ that provided by $R^{(j)}$. Therefore, the same holds in expectation over $(A^{(j)}, B^{(j)})$, which means $\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j+1)})] \geq (1-\epsilon)\mathbf{E}[\widetilde{f}_{A^{(j)},B^{(j)}}(R^{(j)})]$ and the lemma follows. □

4.3 The gap amplification argument

In this section, we develop an inductive argument based on Lemma 4.5 and Lemma 4.7, which proves that a certain notion of density of the distributions at level j increases exponentially in j. By Lemma 4.7, for any X_j constructible at level j there is X_{j+1} constructible at level $j+1$ such that

$$\mathbf{E}[1-(1-\phi(X_{j+1}))^2] \geq (1-\epsilon)\mathbf{E}[\phi((X_j - 10^{-\ell})_+)] - 10^{-\ell}.$$

We want to prove that X_{j+1} is in some sense "significantly larger" than $\frac{1}{2}X_j$. Our main technical lemma formalizing this intuition is the following.

Lemma 4.8. *There are absolute constants $\epsilon, \delta > 0$ such that the following holds for any sufficiently large $\ell \in \mathbb{N}$. If $\mathcal{X}_0, \ldots, \mathcal{X}_\ell$ are collections of random variables in $[0,1]$ such that*

- *there is X_0 in $\mathcal{X}_0$ such that $\mathbf{E}[X_0] \geq c$ for some $c \geq 2^{-\ell}$, and*
- *for every X_j in $\mathcal{X}_j$ and every non-decreasing concave function $\phi : [0,1] \to [0,1]$, there is X_{j+1} in $\mathcal{X}_{j+1}$ such that*

$$\mathbf{E}[1-(1-\phi(X_{j+1}))^2] \geq (1-\epsilon)\mathbf{E}[\phi((X_j - 10^{-\ell})_+)] - 10^{-\ell}$$

then there is a sequence of variables X_j in $\mathcal{X}_j$ and parameters $1 = \alpha_0 \geq \alpha_1 \geq \ldots \alpha_\ell > 0$ such that if we define $\phi_\alpha(t) = \min\left\{\frac{t}{\alpha}, 1\right\}$ then

$$\alpha_j(\mathbf{E}[\phi_{\alpha_j}(X_j)])^{1+\delta} \geq \left(\frac{1+\delta^2}{2}\right)^j c^{1+\delta}.$$

The use of $1+\delta$ in the exponent is crucial here; note that it makes the statement stronger, but this is what makes the inductive proof work. The intuitive meaning of this lemma is as follows: there exist random variables X_j constructible at different levels that, when measured by suitable test functions, decrease roughly as $\left(\frac{1+\delta^2}{2}\right)^j$, rather than $\frac{1}{2^j}$. In terms of the cardinality of the returned sets, this means they increase by a factor of $(1+\delta^2)$ at each level. This gives the exponential amplification that we need.

Proof. The base case $j=0$ holds trivially with $\alpha_0 = 1$ and $\phi_{\alpha_0}(t) = t$. To prove the inductive step, suppose that there is $\alpha_j \in [0,1]$ that satisfies the statement of the lemma for X_j. Let us define $\xi_j = \mathbf{E}[\phi_{\alpha_j}(X_j)]$; then the inductive statement reads

$$\alpha_j \xi_j^{1+\delta} \geq \left(\frac{1+\delta^2}{2}\right)^j c^{1+\delta}. \tag{4}$$

Our goal is to prove that $\alpha_{j+1}\xi_{j+1}^{1+\delta} \geq \frac{1+\delta^2}{2}\alpha_j \xi_j^{1+\delta}$, which implies the inductive statement for $j+1$.

By assumption, for the non-decreasing concave function ϕ_{α_j}, we get

$$\mathbf{E}[1-(1-\phi_{\alpha_j}(X_{j+1}))^2] \geq (1-\epsilon)\mathbf{E}[\phi_{\alpha_j}((X_j - 10^{-\ell})_+)] - 10^{-\ell}.$$

First, we simplify the error terms on the right-hand side. Let us keep in mind that $\epsilon, \delta > 0$ are (small) absolute constants which will be suitably chosen at the end of the proof. Recall that $\phi_{\alpha_j}(t) = \min\{\frac{t}{\alpha_j}, 1\}$. Therefore, $(1-\epsilon)\mathbf{E}[\phi_{\alpha_j}((X_j - 10^{-\ell})_+)] \geq (1-\epsilon)\mathbf{E}[\phi_{\alpha_j}(X_j)] - \frac{1}{\alpha_j 10^\ell} = (1-\epsilon)\xi_j - \frac{1}{\alpha_j 10^\ell}$. Recall the inductive hypothesis (4). Since $\alpha_j, \xi_j \in [0,1]$, and $c \geq 2^{-\ell}$, this means in particular that $\alpha_j \xi_j \geq 2^{-j} c^{1+\delta} \geq 2^{-3\ell}$. Also, $\xi_j \geq 2^{-j} c \geq 2^{-2\ell}$. Hence, we can estimate

$$\mathbf{E}[1-(1-\phi_{\alpha_j}(X_{j+1}))^2] \geq (1-\epsilon)\xi_j - \frac{1}{\alpha_j 10^\ell} - \frac{1}{10^\ell} \geq \left(1 - \epsilon - \frac{2^{3\ell}}{10^\ell} - \frac{2^{2\ell}}{10^\ell}\right)\xi_j \geq (1-2\epsilon)\xi_j \tag{5}$$

for ℓ sufficiently large.

Now we come to the meat of the inductive argument. Instead of the expression $\mathbf{E}[1-(1-\phi_{\alpha_j}(X_{j+1}))^2]$, we would like to estimate $\xi_{j+1} = \mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})]$ for a suitable value of α_{j+1}. The reason why the values of α_j are not specified by the lemma is that their choice depends on the particular distributions of X_j over which we have no control. For example, $\alpha_{j+1} = \frac{1}{2}\alpha_j$ is a natural choice which works for some distributions of X_{j+1} but not always. In the following, we split the analysis into 2 cases.

Case 1: $\Pr[\frac{X_{j+1}}{\alpha_j} > \sqrt{\delta}] > 2\delta\xi_j$.

In this case, X_{j+1} is with non-negligible probability quite large, in the region where $1-(1-X_{j+1}/\alpha_j)^2$ is significantly smaller than $2X_{j+1}/\alpha_j$. In this case, we can gain by making α_{j+1} slightly larger than $\frac{1}{2}\alpha_j$, specifically $\alpha_{j+1} = \frac{1}{2}(1+\delta)\alpha_j$. We obtain:

$$\xi_{j+1} = \mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})] = \mathbf{E}\left[\min\left\{\frac{X_{j+1}}{\alpha_{j+1}}, 1\right\}\right] = \mathbf{E}\left[\min\left\{\frac{2X_{j+1}}{(1+\delta)\alpha_j}, 1\right\}\right] = \frac{1}{1+\delta}\mathbf{E}\left[\min\left\{\frac{2X_{j+1}}{\alpha_j}, 1+\delta\right\}\right].$$

Observe the following: $\min\{\frac{2X_{j+1}}{\alpha_j}, 1+\delta\} \geq \min\{\frac{2X_{j+1}}{\alpha_j}, 1\} \geq 1-(1-\phi_{\alpha_j}(X_{j+1}))^2$ for all $X_{j+1} \geq 0$. Moreover, if $X_{j+1} > \sqrt{\delta}\alpha_j$, we gain an additional δ, because then

$$\min\left\{\frac{2X_{j+1}}{\alpha_j}, 1+\delta\right\} \geq 1 - \left(1 - \min\left\{\frac{X_{j+1}}{\alpha_j}, 1\right\}\right)^2 + \delta = 1 - (1-\phi_{\alpha_j}(X_{j+1}))^2 + \delta$$

(with equality for $X_{j+1} = \sqrt{\delta}\alpha_j$ and $X_{j+1} \geq \alpha_j$; the best way to verify this is to ponder the graph in Figure 4). Therefore,

$$\xi_{j+1} = \frac{1}{1+\delta}\mathbf{E}\left[\min\left\{\frac{2X_{j+1}}{\alpha_j}, 1+\delta\right\}\right] \geq \frac{1}{1+\delta}\mathbf{E}[1-(1-\phi_{\alpha_j}(X_{j+1}))^2] + \frac{\delta}{1+\delta}\Pr[X_{j+1} > \sqrt{\delta}\alpha_j].$$

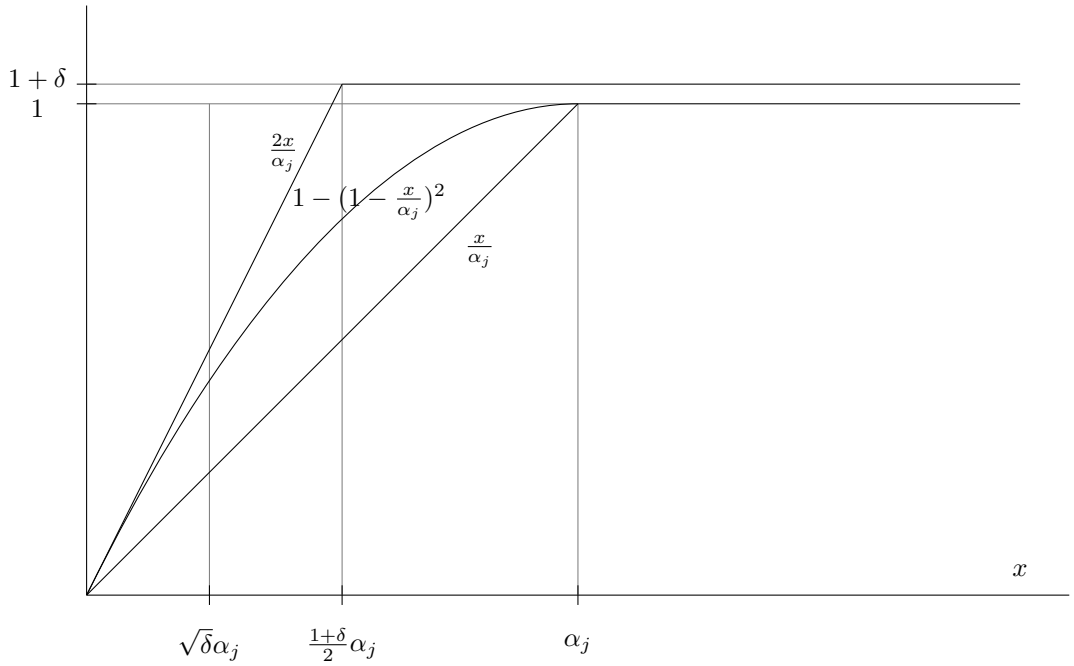

Figure 4: Comparison of the 3 relevant functions for Case 1: Note that for $x \geq \sqrt{\delta}\alpha_j$, the top two functions differ by at least δ; i.e, $\min\{\frac{2x}{\alpha_j}, 1+\delta\} \geq 1-(1-\min\{\frac{x}{\alpha_j}, 1\})^2 + \delta$.

Using (5) and $\Pr[X_{j+1} > \sqrt{\delta}\alpha_j] > 2\delta\xi_j$, we get

$$\xi_{j+1} \geq \frac{1-2\epsilon}{1+\delta}\xi_j + \frac{2\delta^2}{1+\delta}\xi_j = \frac{1+2\delta^2-2\epsilon}{1+\delta}\xi_j.$$

Since $\alpha_{j+1} = \frac{1+\delta}{2}\alpha_j$, we get

$$\alpha_{j+1}\xi_{j+1}^{1+\delta} \geq \frac{1+\delta}{2}\alpha_j \left(\frac{1+2\delta^2-2\epsilon}{1+\delta}\right)^{1+\delta} \xi_j^{1+\delta} = \frac{(1+2\delta^2-2\epsilon)^{1+\delta}}{2(1+\delta)^{\delta}}\alpha_j\xi_j^{1+\delta}.$$

We choose $\epsilon = \delta^4$, so that $(1+2\delta^2-2\epsilon)^{1+\delta} = (1+2\delta^2-2\delta^4)^{1+\delta} \geq 1+2\delta^2+\delta^4$ (it can be verified that this holds for $\delta \in [0, \frac{1}{2}]$). We also use $(1+\delta)^{\delta} \leq 1+\delta^2$ (which holds for $\delta \in [0,1]$). This implies the inductive statement:

$$\alpha_{j+1}\xi_{j+1}^{1+\delta} \geq \frac{1+2\delta^2+\delta^4}{2(1+\delta^2)}\alpha_j\xi_j^{1+\delta} = \frac{1+\delta^2}{2}\alpha_j\xi_j^{1+\delta}.$$

Case 2: $\Pr[\frac{X_{j+1}}{\alpha_j} > \sqrt{\delta}] \leq 2\delta\xi_j$.
In this case, X_{j+1} is almost always very small compared to α_j. Then we can gain by making α_{j+1} much smaller than α_j; we let $\alpha_{j+1} = \sqrt{\delta}\alpha_j$. We have

$$\begin{aligned}
\xi_{j+1} = \mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})] &= \mathbf{E}\left[\min\left\{\frac{X_{j+1}}{\alpha_j\sqrt{\delta}}, 1\right\}\right] = \frac{1}{\sqrt{\delta}}\,\mathbf{E}\left[\min\left\{\frac{X_{j+1}}{\alpha_j}, \sqrt{\delta}\right\}\right] \\
&\geq \frac{1}{\sqrt{\delta}}\left(\mathbf{E}\left[\min\left\{\frac{X_{j+1}}{\alpha_j}, 1\right\}\right] - (1-\sqrt{\delta})\Pr\left[\frac{X_{j+1}}{\alpha_j} > \sqrt{\delta}\right]\right) \\
&\geq \frac{1}{\sqrt{\delta}}\left(\mathbf{E}[\phi_{\alpha_j}(X_{j+1})] - (1-\sqrt{\delta})\cdot 2\delta\xi_j\right)
\end{aligned}$$

An elementary bound together with (5) gives

$$\mathbf{E}[\phi_{\alpha_j}(X_{j+1})] \geq \frac{1}{2}\mathbf{E}[1-(1-\phi_{\alpha_j}(X_{j+1}))^2] \geq \frac{1}{2}(1-2\epsilon)\xi_j.$$

Therefore, using our choice of $\epsilon = \delta^4$,

$$\xi_{j+1} = \mathbf{E}[\phi_{\alpha_{j+1}}(X_{j+1})] \geq \frac{1}{\sqrt{\delta}}\left(\frac{1}{2}(1-2\epsilon)\xi_j - 2(1-\sqrt{\delta})\delta\xi_j\right) = \frac{1-2\delta^4-4\delta+4\delta^{3/2}}{2\sqrt{\delta}}\xi_j \geq \frac{1-4\delta+2\delta^{3/2}}{2\sqrt{\delta}}\xi_j.$$

From here, using $\alpha_{j+1} = \sqrt{\delta}\alpha_j$ and $(1-4\delta+2\delta^{3/2})^{1+\delta} \geq 1-4\delta$ (which holds for any $\delta \in [0, \frac{1}{4}]$),

$$\alpha_{j+1}\xi_{j+1}^{1+\delta} \geq \sqrt{\delta}\alpha_j\left(\frac{1-4\delta+2\delta^{3/2}}{2\sqrt{\delta}}\right)^{1+\delta}\xi_j^{1+\delta} \geq \frac{1-4\delta}{2^{1+\delta}\delta^{\delta/2}}\alpha_j\xi_j^{1+\delta}.$$

We choose $\delta = e^{-10}$ so that $\delta^{\delta/2} = e^{-5\delta}$. Then,

$$\alpha_{j+1}\xi_{j+1}^{1+\delta} \geq \frac{1-4\delta}{2^{1+\delta}}e^{5\delta}\alpha_j\xi_j^{1+\delta} \geq \frac{1+\delta^2}{2}\alpha_j\xi_j^{1+\delta}$$

which finishes the inductive step. □

Putting together Lemma 4.8 and the cardinality bound which applies to every feasible solution, we complete our hardness result for combinatorial public projects.

Proof of Theorem 4.2. Let $\epsilon > 0$ and $\delta > 0$ be the constants provided by Lemma 4.8. Let $n = 2^\ell$ and $m = 400^\ell$. Suppose there is a mechanism for the problem $\max\{f(S) : |S| \leq m/n\}$ that maximizes $(1-\epsilon)$-approximately over a distributional range $\mathcal{R}$ and provides a c-approximation, where $c \geq 1/n$. By Lemma 4.7 and Lemma 4.5, there are collections of random variables $\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_\ell$ constructible at the respective levels by $\mathcal{R}$, satisfying the conditions of Lemma 4.8. Hence, by Lemma 4.8 for $j = \ell$, there is X_ℓ constructible by $\mathcal{R}$ at level ℓ such that

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

Recall that $\phi_{\alpha_\ell}(t) = \min\{\frac{t}{\alpha_\ell}, 1\}$. Therefore, we have

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

We have $X_\ell = \frac{|R|}{|M|}$ where R is a random set sampled according to some distribution in the range $\mathcal{R}$. All distributions in the range must be feasible in expectation, otherwise the mechanism cannot possibly maximize over them and return a feasible solution. Therefore, $\mathbf{E}[X_\ell] \leq \frac{1}{n}$ which implies that

$$c \leq (1+\delta^2)^{-\frac{\ell}{1+\delta}} < 2^{-\delta^2\ell} = n^{-\delta^2}.$$

Therefore, there is no $(1-\epsilon)$-approximately MIDR mechanism providing an $n^{-\delta^2}$-approximation in the objective function. Also, we have $m = 400^\ell = \text{poly}(n)$, so the approximation cannot be better than $m^{-\gamma}$ for some constant $\gamma > 0$. The only bound we have used on the mechanism was that the number of value queries is polynomial in n, or equivalently polynomial in m. □

5 Hardness for combinatorial auctions

In this section, we present 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 in the cases of deterministic truthful and randomized universally truthful mechanisms [16]. Also, these results do not rule out $(1-\epsilon)$-approximately truthful mechanisms.

We remark that there is still the possibility of a truthful-in-expectation mechanism in the "lottery-value" oracle model which was introduced in [20], or perhaps some other oracle model. In the lottery-value oracle model, 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 #P-hard to compute for valuation functions of interest, such as matroid rank functions (see Appendix C), this is a significant 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 t.i.e. 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 t.i.e. 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 the case of 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 [10]). 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 [16]. The *direct hardness* approach of [16] 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 formal proof of Theorem 5.1.

5.1 The basic random instance

We choose a parameter $\ell \geq 1$ and construct instances with $|N| = n = 2^\ell$ players and $|M| = m = 400^\ell$ items. We define "polar valuations" as in [16].

Definition 5.2. *Given a set of items $A \subset M$ and a parameter $\omega > 0$, the polar valuation v_A^* associated with A is defined by*

$$v_A^*(S) = |A \cap S| + \omega |S \setminus A|.$$

Our "basic instance" is an instance where each player has a polar valuation associated with a random set of size m/n.

Definition 5.3. *In the basic instance, player i has valuation $v_i^* = v_{A_i^{(0)}}^*$ where $A_i^{(0)}$ is a uniformly random set of size m/n, chosen independently for each player.*

Next, we prove that for some player, his allocation overlaps significantly with his desired set.

Lemma 5.4. *For any c-approximation mechanism applied to the random basic instance, there is a player i and sets $A_j^{(0)}, j \neq i$, such that conditioned on the desired sets for players $j \neq i$ being $A_j^{(0)}$, player i gets allocated a random set $R_i^{(0)}$ such that*

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

Proof. First, let us estimate the optimal social welfare that the basic instance admits in expectation. Given $(A_1^{(0)}, \ldots, A_n^{(0)})$, each item in $\bigcup_{i=1}^n A_i^{(0)}$ can be allocated to some player so that it brings value 1. We ignore the remaining items. Observe that a fixed item j appears in each $A_i^{(0)}$ independently with probability $1/n$, therefore $Pr[j \in \bigcup_{i=1}^n A_i^{(0)}] = 1 - (1-1/n)^n \geq 1 - 1/e > 1/2$. Hence,

$$\mathbf{E}[OPT] \geq \mathbf{E}[|\bigcup_{i=1}^n A_i^{(0)}|] > \frac{m}{2}.$$

We remind the reader that the expectation is over the random choices of $(A_1^{(0)}, \ldots, A_n^{(0)})$. A c-approximate mechanism should provide at least $c \cdot OPT$ in expectation for every particular instance. Hence also in expectation over the random choice of $(A_1^{(0)}, \ldots, A_n^{(0)})$. If $(R_1, \ldots, R_n)$ is the allocation provided by the mechanism, this means

$$\sum_{i=1}^n \mathbf{E}[|R_i \cap A_i^{(0)}| + \omega|R_i \setminus A_i^{(0)}|] \geq c \cdot OPT > \frac{cm}{2}.$$

Since each of $A_1^{(0)}, \ldots, A_n^{(0)}$ has size m/n and the sizes of $R_1, \ldots, R_n$ add up to at most m, we can write

$$\sum_{i=1}^n \mathbf{E}[|R_i \cap A_i^{(0)}| + \omega|R_i \setminus A_i^{(0)}|] > \frac{cm}{2} \geq \frac{c}{4}\sum_{i=1}^n \mathbf{E}[|R_i| + |A_i^{(0)}|] \geq \frac{c}{4}\sum_{i=1}^n \mathbf{E}[|R_i \cup A_i^{(0)}|].$$

By an averaging argument, there must be i such that

$$\mathbf{E}[|R_i \cap A_i^{(0)}| + \omega|R_i \setminus A_i^{(0)}|] > \frac{c}{4}\mathbf{E}[|R_i \cup A_i^{(0)}|]$$

and therefore

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

This holds in expectation over the choices of $(A_j^{(0)} : j \neq i)$, and again by an averaging argument it also holds conditioned on some particular choice of $(A_j^{(0)} : j \neq i)$. We call the random set allocated to player i under this conditioning $R_i^{(0)}$. □

5.2 Setup for higher-level valuations

In the following, the valuations of all players except i are fixed to be $v_j^* = v^*_{A_j^{(0)}}$ for some choice of sets $(A_j^{(0)} : j \neq i)$. Now we will vary the valuation of player i in order to be able to apply the symmetry gap argument as before. Since we work only with player i, we call him the "special player" and we drop the index i in the following.

Recall Definition 4.3, the definition of a random bisection sequence. We will use the same concept here, where at level j we have a random set $A^{(j)}$ partitioned randomly into $A^{(j-1)} \cup B^{(j-1)}$. These sets have sizes $|A^{(j-1)}| = |B^{(j-1)}| = 2^{j-1-\ell}m$. We will use valuation functions associated with each pair of sets. We denote these valuation functions by $v_{A^{(j-1)},B^{(j-1)}}$ at level j. In particular, this valuation function depends only on the elements of $A^{(j)} = A^{(j-1)} \cup B^{(j-1)}$. In other words, $A^{(j-1)}$ and $B^{(j-1)}$ are the desired sets of items at level j.

Distribution menu. For each particular choice of a valuation function $v_{A^{(j)},B^{(j)}}$ at a certain level, the mechanism needs to produce a distribution over item sets for the special player, along with a certain price. Recall that due to the definition of (approximate) truthfulness in expectation, this choice should give (approximately) the optimal utility for the special player among all possible choices given the other valuations v^*_{-i}. After fixing a set of valuation functions $v_{A^{(j)},B^{(j)}}$ for each pair $(A^{(j)}, B^{(j)})$, the output distribution will depend only on $(A^{(j)}, B^{(j)})$ and hence we denote the respective random set by $R(A^{(j)}, B^{(j)})$; we also denote the associated price by $P(A^{(j)}, B^{(j)})$. Thus the mechanisms assigns distributions over sets and prices to all pairs of sets. As before, we distill the important information from the distribution into a random variable X_j. There is some additional information now expressed by the price; we associate the price with a separate random variable P_j. The possible choices of distributions for (X_j, P_j) are what we call a *distribution menu* at level j.

Definition 5.5. *Given a mechanism and a special player with other valuations fixed, the "distribution menu at level j", $\mathcal{M}_j$, is the set of all probability distributions of a pair of variables (X_j, P_j) that arise as follows: There exist valuations $v_{A^{(j-1)},B^{(j-1)}}$ such that when declaring $v_{A^{(j-1)},B^{(j-1)}}$, the special player receives a random set $R(A^{(j-1)}, B^{(j-1)})$ at a price $P(A^{(j-1)}, B^{(j-1)})$. Then, for $A^{(j)} = A^{(j-1)} \cup B^{(j-1)}$ chosen as the $(j-1)$-th level of a random bisection sequence, i.e. a random pair of disjoint sets of size $2^{j-1-\ell}m$, we have*

$$X_j = \frac{|A^{(j)} \cap R(A^{(j-1)}, B^{(j-1)})|}{|A^{(j)}|},$$

$$P_j = P(A^{(j-1)}, B^{(j-1)}).$$

In other words, X_j encodes the (random) fraction of the relevant items that the special player receives at level j, and P_j is the respective (random) price. Note that there are two sources of randomness in (X_j, P_j): one is the random choice of $(A^{(j-1)}, B^{(j-1)})$, and one arises from the randomness of the mechanism for fixed $(A^{(j-1)}, B^{(j-1)})$.

Closure of a distribution menu. Furthermore, it will be convenient to make the menu closed and convex as follows.

Definition 5.6. *We define $\overline{\mathcal{M}_j}$, the closure of the distribution menu at level j, to be the topological closure of the set of all convex combinations of distributions from the menu $\mathcal{M}_j$.*

I.e., we take the convex hull of the menu and then its topological closure. We emphasize that the convex hull is generated by averaging distributions, and not the values of (X_j, P_j). In other words, a distribution of (X'_j, P'_j) is in $\overline{\mathcal{M}_j}$ if its distribution can be approximated arbitrarily closely by some convex combination of distributions in $\mathcal{M}_j$. It is important that we keep all the randomness present in (X_j, P_j) and do not take expectations until the end.

5.3 Symmetry gap revisited

Recall Lemma 4.6 which was proved using the symmetry gap argument and played an important role in our proof for the CPP problem. We still want to use this lemma; however, the difference now is the presence of prices. In order to deal with prices, we need to introduce a parameter λ which acts as a conversion factor between values and prices. For that purpose, we prove the following slight variation of Lemma 4.6.

Lemma 5.7. *Let $\phi : [0,1] \to [0,1]$ be a non-decreasing concave function and $\lambda \geq 0$ any constant. Let $A^{(j+1)}$ be a fixed set of size $2^{j+1-\ell}m \geq m/n$ and $(A^{(j)}, B^{(j)})$ a random partition of $A^{(j+1)}$ into two sets of equal size. Then there is a monotone submodular function $\widetilde{v}_{A^{(j)},B^{(j)}}$ for each partition $(A^{(j)}, B^{(j)})$ such that*

- *For any distribution of a random set $R^{(j)}$ (possibly correlated with $A^{(j)}$) and the associated random variable $X_j = \frac{|R^{(j)} \cap A^{(j)}|}{|A^{(j)}|}$, we have*

$$\mathbf{E}[\widetilde{v}_{A^{(j)},B^{(j)}}(R^{(j)})] \geq \lambda \mathbf{E}[\phi((X_j - nm^{-1/2})_+)].$$

- *Any mechanism that uses poly(n) value queries, when applied to the random input $\widetilde{v}_{A^{(j)},B^{(j)}}$ will return a random set $R^{(j+1)}$ such that for the random variable $X_{j+1} = \frac{|R^{(j+1)} \cap A^{(j+1)}|}{|A^{(j+1)}|}$,*

$$\mathbf{E}[\widetilde{v}_{A^{(j)},B^{(j)}}(R^{(j+1)})] \leq \lambda \mathbf{E}[1 - (1 - \phi(X_{j+1}))^2 + e^{-\Omega(n)}].$$

The expectations are over both $(A^{(j)}, B^{(j)})$ and $R^{(j)}$ or $R^{(j+1)}$ respectively.

Proof. The proof is easily obtained from the proof of Lemma 4.6. The only difference is the scaling by $\lambda \geq 0$. (For $\lambda = 0$ the statement is trivial.) Given $\phi : [0,1] \to [0,1]$ and $\lambda \geq 0$, we take the function $\widetilde{f}_{A^{(j)},B^{(j)}}$ provided by Lemma 4.6 and scale it by λ:

$$\widetilde{v}_{A^{(j)},B^{(j)}}(S) = \lambda \widetilde{f}_{A^{(j)},B^{(j)}}(S).$$

Randomizing over $(A^{(j)}, B^{(j)})$, the same proof shows that any mechanism will return a random set $R^{(j+1)}$ with high probability balanced with respect to $(A^{(j)}, B^{(j)})$. Hence we obtain the same bounds as in Lemma 4.6 with the right-hand side scaled by λ. □

Applying the assumption of approximate truthfulness, we obtain the following.

Lemma 5.8. *Consider a $(1-\epsilon)$-approximately truthful-in-expectation mechanism for combinatorial auctions with $n = 2^\ell$ players and $m = 400^\ell$ items. Let $\phi : [0,1] \to [0,1]$ be a non-decreasing concave function. If (X_j, P_j) has a distribution in the closure of the level-j menu $\overline{\mathcal{M}_j}$, then for any $\lambda', \lambda'' \geq 0$ there is (X_{j+1}, P_{j+1}) with a distribution in the closure of the level-$(j+1)$ menu $\overline{\mathcal{M}_{j+1}}$ such that*

$$\lambda' \mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] - \lambda'' \mathbf{E}[P_{j+1}] \geq \lambda' \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell}] - \lambda'' \mathbf{E}[P_j].$$

Proof. First let us assume that $\lambda'' > 0$ and set $\lambda = \lambda'/\lambda''$. If (X_j, P_j) is on the menu $\mathcal{M}_j$, it means that the mechanism under certain valuations depending on the (random) set $A^{(j)}$ allocates to the special player a random set $R^{(j)}$ (at some price P_j) such that $X_j = \frac{|R^{(j)} \cap A^{(j)}|}{|A^{(j)}|}$. Given ϕ, by Lemma 5.7 there are valuation functions $\widetilde{v}_{A^{(j)},B^{(j)}}$ such that

$$\mathbf{E}[\widetilde{v}_{A^{(j)},B^{(j)}}(R^{(j)})] \geq \lambda \mathbf{E}[\phi((X_j - nm^{-1/2})_+)]$$

and on the other hand, the mechanism executed on this random input allocates a random set $R^{(j+1)}$ such that with $X_{j+1} = \frac{|R^{(j+1)} \cap A^{(j+1)}|}{|A^{(j+1)}|}$,

$$\mathbf{E}[\widetilde{v}_{A^{(j)},B^{(j)}}(R^{(j+1)})] \leq \lambda \mathbf{E}[1 - (1 - \phi(X_{j+1}))^2 + e^{-\Omega(n)}].$$

Let us assume that the mechanism allocates this distribution at a price P_{j+1}. Due to the assumption of $(1-\epsilon)$-truthfulness, the utility provided by the mechanism must be approximately maximized for the true valuation. Hence, we must have

$$\lambda \mathbf{E}[1 - (1 - \phi(X_{j+1}))^2 + e^{-\Omega(n)}] - \mathbf{E}[P_{j+1}] \geq (1-\epsilon)\lambda \mathbf{E}[\phi((X_j - nm^{-1/2})_+)] - \mathbf{E}[P_j]. \tag{6}$$

Given our parameters $m = 400^\ell, n = 2^\ell$, we have $nm^{-1/2} = 10^{-\ell}$ and $e^{-\Omega(n)} = e^{-\Omega(2^\ell)} << 10^{-\ell}$, therefore

$$\lambda \mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] - \mathbf{E}[P_{j+1}] \geq \lambda \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell}] - \mathbf{E}[P_j].$$

Since this inequality is preserved under convex combinations and limits of the distributions of (X_j, P_j) and (X_{j+1}, P_{j+1}) (the non-linearity of ϕ is irrelevant here!), the same holds for the closures $\overline{\mathcal{M}_j}, \overline{\mathcal{M}_{j+1}}$: For any (X_j, P_j) with a distribution in $\overline{\mathcal{M}_j}$ and $\lambda > 0$, there exists (X_{j+1}, P_{j+1}) with a distribution in $\overline{\mathcal{M}_{j+1}}$ such that (6) holds. This proves the statement of the lemma when $\lambda'' > 0$.

When $\lambda'' = 0$, the statement claims that given $(X_j, P_j) \in \overline{\mathcal{M}_j}$, there is $(X_{j+1}, P_{j+1}) \in \overline{\mathcal{M}_{j+1}}$, such that $\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] \geq \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell}]$, without regard to prices. This can be

obtained from the previous discussion as follows. Let us assume that p^* is an absolute lower bound on the expected price $\mathbf{E}[P_{j+1}]$. (If the mechanism possibly pays arbitrarily large amounts on the menu of the special player, then given a zero valuation it cannot maximize utility over the menu.) Given (X_j, P_j) in $\overline{\mathcal{M}_j}$, let $\lambda = 10^{\ell+1}(\mathbf{E}[P_j] - p^*)$; there must be a pair (X_{j+1}, P_{j+1}) in $\overline{\mathcal{M}_{j+1}}$ satisfying (6). Using the (still very crude) estimate $e^{-\Omega(n)} << 10^{-\ell-1}$, (6) implies

$$\begin{aligned}
\mathbf{E}[1-(1-\phi(X_{j+1}))^2] &\geq \mathbf{E}[(1-\epsilon)\phi((X_j - nm^{-1/2})_+) - e^{-\Omega(n)}] - \frac{\mathbf{E}[P_j] - p^*}{\lambda} \\
&\geq \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell-1}] - 10^{-\ell-1} \\
&\geq \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell}].
\end{aligned}$$

□

5.4 The convex separation argument

Next, we use a geometric argument, essentially Farkas' lemma in 2 dimensions, which shows that the bounds for varying multipliers $\lambda', \lambda'' \geq 0$ allow us to obtain separate bounds on value and price.

Lemma 5.9. *Consider a $(1-\epsilon)$-approximately truthful-in-expectation mechanism for combinatorial auctions with $n = 2^\ell$ players and $m = 400^\ell$ items. Let $\phi : [0,1] \to [0,1]$ be a non-decreasing concave function. If (X_j, P_j) has a distribution in the closure of the level-j menu $\overline{\mathcal{M}_j}$, then there is (X_{j+1}, P_{j+1}) with a distribution in the closure of the level-$(j+1)$ menu $\overline{\mathcal{M}_{j+1}}$ such that*

$$\mathbf{E}[P_{j+1}] \leq \mathbf{E}[P_j]$$

and

$$\mathbf{E}[1-(1-\phi(X_{j+1}))^2] \geq \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell}].$$

Proof. Denote $q_j = \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell}]$ and $q_{j+1} = \mathbf{E}[1-(1-\phi(X_{j+1}))^2]$. Set also $p_j = \mathbf{E}[P_j]$ and $p_{j+1} = \mathbf{E}[P_{j+1}]$. By Lemma 5.8, for any $(X_j, P_j) \in \overline{\mathcal{M}_j}$ and any $\lambda', \lambda'' \geq 0$, there is $(X_{j+1}, P_{j+1}) \in \overline{\mathcal{M}_{j+1}}$ such that $\lambda' q_{j+1} - \lambda'' p_{j+1} \geq \lambda' q_j - \lambda'' p_j$. Using these transformations, let us map the set $\overline{\mathcal{M}_j}$ to

$$\mathcal{Q}_j = \{(q_j, p_j) : (X_j, P_j) \in \overline{\mathcal{M}_j}\}.$$

and map $\overline{\mathcal{M}_{j+1}}$ to

$$\mathcal{Q}_{j+1} = \{(q_{j+1}, p_{j+1}) : (X_{j+1}, P_{j+1}) \in \overline{\mathcal{M}_{j+1}}\}.$$

Both $\mathcal{Q}_j$ and $\mathcal{Q}_{j+1}$ are closed convex sets, because they are the images of the closed convex sets $\overline{\mathcal{M}_j}, \overline{\mathcal{M}_{j+1}}$ under a linear map (the map being the expectation of a certain function over a distribution; this is linear as a function of the distribution even though the function is non-linear).

By this transformation, we have reduced the proof to a geometric question in the plane (see Figure 5): Given (q_j, p_j), assume that for any $\lambda', \lambda'' \geq 0$, there is $(q_{j+1}, p_{j+1}) \in \mathcal{Q}_{j+1}$ such that $\lambda' q_{j+1} - \lambda'' p_{j+1} \geq \lambda' q_j - \lambda'' p_j$. Is it possible that there is no point $(q_{j+1}, p_{j+1}) \in \mathcal{Q}_{j+1}$ such that $q_{j+1} \geq q_j$ and $p_{j+1} \leq p_j$?

Suppose that there is no such point in $\mathcal{Q}_{j+1}$. This means that $\mathcal{Q}_{j+1}$ and $\{(q,p) : q \geq q_j, p \leq p_j\}$ are disjoint. Since these are closed convex sets, they can be separated by a line. This line cannot have a negative slope, otherwise it would intersect the quadrant $\{(q,p) : q \geq q_j, p \leq p_j\}$. Such a separating line gives $\lambda', \lambda'' \geq 0$ such that $\lambda' q_{j+1} - \lambda'' p_{j+1} < \lambda' q_j - \lambda'' p_j$ for all $(q_{j+1}, p_{j+1}) \in \mathcal{Q}_{j+1}$. However, this contradicts the assumption above. Hence there is a point $(q_{j+1}, p_{j+1}) \in \mathcal{Q}_{j+1}$ such that $q_{j+1} \geq q_j$ and $p_{j+1} \leq p_j$. □

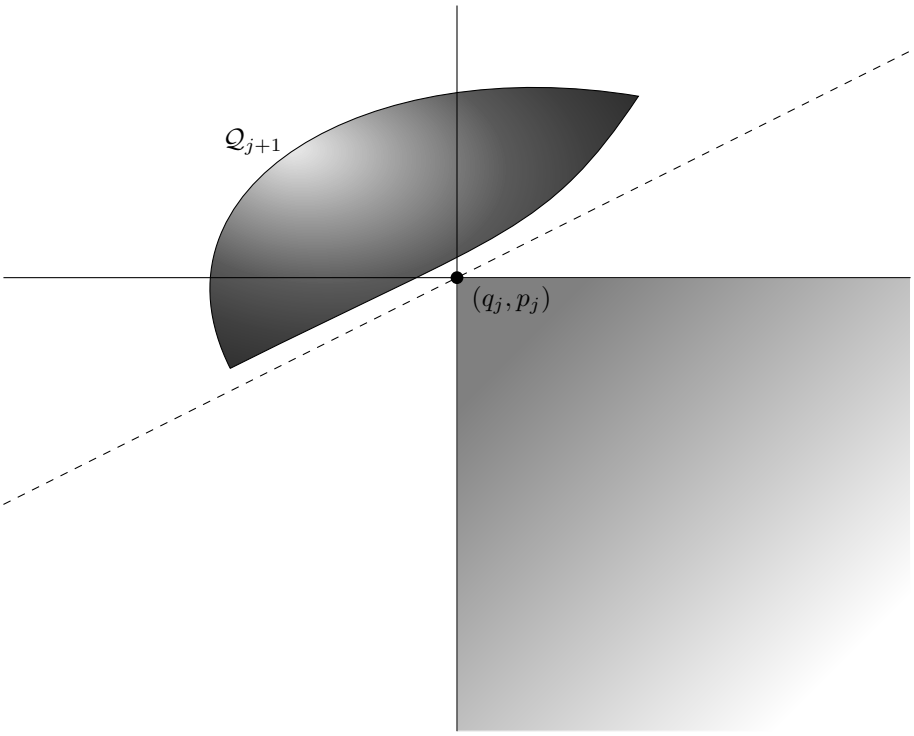

Figure 5: The convex separation argument.

5.5 Putting it all together

In this section, we finish the proof of our main hardness result for combinatorial auctions.

Proof of Theorem 5.1. Let $\epsilon > 0$ and $\delta > 0$ be the constants provided by Lemma 4.8. Let the number of players be $n = 2^\ell$ and the number of items $m = 400^\ell$. Suppose there is a $(1-\epsilon)$-approximately truthful-in-expectation mechanism that provides a c-approximation in social welfare, where $c = 1/n^\gamma = 2^{-\gamma\ell}$ for some constant $\gamma > 0$.

Consider the basic instance (Section 5.1). Choose a special player and fix the remaining valuations, based on Lemma 5.4. Let $R^{(0)}$ be the random set allocated to the special player, P_0 the respective price, $A^{(0)}$ his desired set, $X_0 = \frac{|R^{(0)} \cap A^{(0)}|}{|A^{(0)}|}$, $c_0 = \mathbf{E}[X_0]$ and $p_0 = \mathbf{E}[P_0]$. Lemma 5.4 implies $c_0 \geq c/4 - \omega$. We set $\omega = c/8$. Then $c_0 \geq c/8 = 1/(8n^\gamma) = 2^{-\gamma\ell - 3}$.

Now consider the distribution menus at different levels and their closures $\overline{\mathcal{M}_j}$ (Section 5.2). Let us define $\mathcal{X}_j$ to be the collection of random variables X_j such that (X_j, P_j) is in $\overline{\mathcal{M}_j}$ for some price P_j such that $\mathbf{E}[P_j] \leq p_0$. As discussed above, we have X_0 in $\mathcal{X}_0$ such that $\mathbf{E}[X_0] = c_0 \geq 2^{-\gamma\ell-3} \geq 2^{-\ell}$ for ℓ sufficiently large. Also, Lemma 5.9 says that for any X_j in $\mathcal{X}_j$ and any non-decreasing concave $\phi : [0,1] \to [0,1]$, we have X_{j+1} in $\mathcal{X}_{j+1}$ such that

$$\mathbf{E}[1 - (1 - \phi(X_{j+1}))^2] \geq \mathbf{E}[(1-\epsilon)\phi((X_j - 10^{-\ell})_+) - 10^{-\ell}].$$

In other words, the collections $\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_\ell$ satisfy the assumptions of Lemma 4.8. Hence, by Lemma 4.8 for $j = \ell$, there is X_ℓ in $\mathcal{X}_\ell$ and $\alpha_\ell \in [0,1]$ such that

$$\alpha_\ell (\mathbf{E}[\phi_{\alpha_\ell}(X_\ell)])^{1+\delta} \geq \left(\frac{1+\delta^2}{2}\right)^\ell c_0^{1+\delta}.$$

Recall that $\phi_{\alpha_\ell}(t) = \min\{\frac{t}{\alpha_\ell}, 1\}$. Therefore, we have

$$\mathbf{E}[X_\ell] \geq \alpha_\ell \mathbf{E}[\phi_{\alpha_\ell}(X_\ell)] \geq \alpha_\ell (\mathbf{E}[\phi_{\alpha_\ell}(X_\ell)])^{1+\delta} \geq \left(\frac{1+\delta^2}{2}\right)^\ell c_0^{1+\delta} \geq 2^{\delta^2\ell-\ell} c_0^{1+\delta}.$$

Since X_ℓ is in $\mathcal{X}_\ell$, the respective price is bounded by $\mathbf{E}[P_\ell] \leq p_0$. Now consider an expression related to the utility the special player would derive in the basic instance:

$$\mathbf{E}[(1-\epsilon)\omega m X_\ell - P_\ell] \geq (1-\epsilon)\omega m 2^{\delta^2\ell-\ell} c_0^{1+\delta} - p_0.$$

The distribution of (X_ℓ, P_ℓ) might not be on the actual menu $\mathcal{M}_\ell$ of the special player; however, since it is in the closure of its convex hull, there exists a pair $(\widetilde{X}_\ell, \widetilde{P}_\ell)$ with a distribution in $\mathcal{M}_\ell$ such that

$$\mathbf{E}[(1-\epsilon)\omega m \widetilde{X}_\ell - \widetilde{P}_\ell] > (1-2\epsilon)\omega m 2^{\delta^2\ell-\ell} c_0^{1+\delta} - p_0.$$

(If not, we get a contradiction since if the reverse inequality holds for $\mathcal{M}_\ell$, it holds also for the closure $\overline{M}_\ell$.) The random variable $\widetilde{X}_\ell$ represents a random set $\widetilde{R}^{(\ell)}$, possibly allocated to the special player at price $\widetilde{P}_\ell$: $\widetilde{X}_\ell = \frac{|\widetilde{R}^{(\ell)}|}{|A^{(\ell)}|} = \frac{1}{m}|\widetilde{R}^{(\ell)}|$.

Now let us go back to the basic instance. Considering that the valuation of the special player in the basic instance satisfies $v_i^*(S) \geq \omega|S|$, we obtain

$$\mathbf{E}[(1-\epsilon)v_i^*(\widetilde{R}^{(\ell)}) - \widetilde{P}_\ell] \geq \mathbf{E}[(1-\epsilon)\omega m \widetilde{X}_\ell - \widetilde{P}_\ell] > (1-2\epsilon)\omega m 2^{\delta^2\ell-\ell} c_0^{1+\delta} - p_0.$$

Using $c_0 \geq c/8 = \omega = 2^{-\gamma\ell-3}$, we get

$$\mathbf{E}[(1-\epsilon)v_i^*(\widetilde{R}^{(\ell)}) - \widetilde{P}_\ell] > (1-2\epsilon)\left(\frac{c}{8}\right)^{1+\delta} 2^{\delta^2\ell-\ell} m c_0 - p_0 = (1-2\epsilon)2^{\delta^2\ell-(1+\delta)(\gamma\ell+3)-\ell} m c_0 - p_0. \quad (7)$$

On the other hand, the set $R^{(0)}$ actually allocated under declared valuation v_i^* gives

$$\mathbf{E}[v_i^*(R^{(0)})] = \mathbf{E}[|R^{(0)} \cap A^{(0)}|] + \omega\mathbf{E}[|R^{(0)} \setminus A^{(0)}|] \leq \frac{m}{n}\mathbf{E}[X_0] + \omega\mathbf{E}[|R^{(0)}|] \leq \frac{2m}{n}\mathbf{E}[X_0]$$

using again Lemma 5.4 to say that $\frac{m}{n}\mathbf{E}[X_0] = \mathbf{E}[|R^{(0)} \cap A^{(0)}|] \geq (c/4-\omega)\mathbf{E}[|R^{(0)}|] = \omega\mathbf{E}[|R^{(0)}|]$. Therefore, since $\mathbf{E}[X_0] = c_0$ and $\mathbf{E}[P_0] = p_0$,

$$\mathbf{E}[v_i^*(R^{(0)}) - P_0] \leq \frac{2m}{n}\mathbf{E}[X_0] - \mathbf{E}[P_0] = 2^{1-\ell} m c_0 - p_0. \quad (8)$$

Since $\widetilde{R}^{(\ell)}$ is a random set the special player could receive at price $\widetilde{P}_\ell$ if he had declared a suitable valuation, $(1-\epsilon)$-approximate truthfulness implies that

$$\mathbf{E}[v_i^*(R^{(0)}) - P_0] \geq \mathbf{E}[(1-\epsilon)v_i^*(\widetilde{R}^{(\ell)}) - \widetilde{P}_\ell].$$

Considering (7) and (8), this implies

$$2^{1-\ell} > (1-2\epsilon)2^{\delta^2\ell-(1+\delta)(\gamma\ell+3)-\ell}.$$

We conclude that $\gamma \geq \frac{\delta^2}{1+\delta}$, otherwise we get a contradiction for a large enough ℓ. □

References

[1] N. Alon and J.H. Spencer. The probabilistic method (2nd edition). Wiley Interscience, 2000.

[2] L. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, 2006.

[3] L. Blumrosen and N. Nisan. Combinatorial auctions (a survey). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[4] L. Blumrosen and N. Nisan. Informational limitations of ascending combinatorial auctions. *Journal of Economic Theory*, 145(3):1203–1223, 2010.

[5] 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* 21^{th} *ACM-SIAM SODA*, 518–536, 2010.

[6] D. Buchfuhrer, M. Schapira and Y. Singer. Computation and incentives in combinatorial public projects. *Proc. of* 11^{th} *ACM EC*, 33–42, 2010.

[7] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. *Proc. of* 12^{th} *IPCO*, 182–196, 2007.

[8] D. Chakrabarty and G. Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. *Proc. of* 49^{th} *IEEE FOCS*, 687–696, 2008.

[9] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press., 2006.

[10] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. *Proc. of APPROX 2007*, 89–103.

[11] S. Dobzinski and S. Dughmi, On the power of randomization in algorithmic mechanism design *Proc. of* 50^{th} *IEEE FOCS*, 505–514, 2009.

[12] S. Dobzinski and N. Nisan. Limitations of VCG-based mechanisms. *Proc. of* 39^{th} *ACM STOC*, 338–344, 2007.

[13] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proc.* 37^{th} *ACM STOC*, pages 644–652, 2006.

[14] S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. *Proc. of* 17^{th} *SODA*, 1064–1073, 2006.

[15] S. Dobzinski, H. Fu, and R. Kleinberg. Truthfulness via proxies. *arXiv:1011.3232*, 2010.

[16] S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. *Proc. of* 43^{rd} *ACM STOC*, 2011.

[17] S. Dobzinski and J. Vondrák: From query complexity to computational complexity. *Proc. of* 44^{th} *ACM STOC*, 1107–1116, 2012.

[18] S. Dobzinski and J. Vondrák. The computational complexity of truthfulness in combinatorial auctions. *Proc. of* 13^{th} *ACM EC*, 405–422, 2012.

[19] S. Dughmi and T. Roughgarden. Black box randomized reductions in algorithmic mechanism design. *Proc. of* 51^{st} *IEEE FOCS*, 775–784, 2010.

[20] S. Dughmi, T. Roughgarden and Q. Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. *Proc. of* 43^{rd} *ACM STOC*, 149–158, 2011.

[21] S. Dughmi, T. Roughgarden, J. Vondrák and Q. Yan. An approximately truthful-in-expectation mechanism for combinatorial auctions using value queries. Manuscript, arXiv:1109.1053, 2011.

[22] S. Dughmi. A truthful randomized mechanism for combinatorial public projects via convex optimization. In 12^{th} *ACM EC 2011*.

[23] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.

[24] U. Feige. On maximizing welfare where the utility functions are subadditive. In *Proceedings of the 37th ACM Symposium on Theory of Computing (STOC)*, pages 122–142, 2006.

[25] U. Feige, V. Mirrokni and J. Vondrák. Maximizing a non-monotone submodular function. *Proc. of* 48^{th} *IEEE FOCS*, 461–471, 2007.

[26] U. Feige and J. Vondrák. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. *Proc. of* 47^{th} *IEEE FOCS*, 667–676, 2006.

[27] S. Khot, R. Lipton, E. Markakis and A. Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica* 52:1, 3–18, 2008.

[28] A. Kulik, H. Shachnai and T. Tamir. Maximizing submodular set functions subject to multiple linear constraints. *Proc. of* 20^{th} *ACM-SIAM SODA*, 545–554, 2009.

[29] R. Lavi, A. Mu'alem and N. Nisan. Towards a characterization of truthful combinatorial auctions. *Proc. of* 44^{th} *IEEE FOCS*, 574–583, 2003.

[30] R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. *Proc. of* 46^{th} *IEEE FOCS*, 595–604, 2005.

[31] B. Lehmann, D. Lehmann and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Proc. of* 3^{rd} *ACM EC*, 18–28, 2001.

[32] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[33] V. Mirrokni, M. Schapira and J. Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. *Proc. of* 9^{th} *ACM EC*, 70–77, 2008.

[34] 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.

[35] 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.

[36] S. Oveis Gharan and J. Vondrák. Submodular maximization by simulated annealing. *Proc. of* 22^{nd} *ACM-SIAM SODA*, 1098-1117, 2011.

[37] J. Oxley. Matroid Theory, Cambridge University Press, 1992.

[38] C. Papadimitriou, M. Schapira and Y. Singer. On the hardness of being truthful. *Proc. of* 49^{th} *IEEE FOCS*, 250–259, 2008.

[39] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. *Proc. of* 40^{th} *ACM STOC*, 67–74, 2008.

[40] J. Vondrák. Symmetry and approximability of submodular maximization problems. *Proc. of IEEE FOCS*, 251–270, 2009.

A Transforming an approximately truthful-in-expectation mechanism for CPP into approximately maximal-in-distribution range

In this section, we prove 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 $\mathbb{R}^\Omega$, 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 $\mathbb{R}_+^\Omega$ — specifically, each $v \in \mathcal{V}$ is a vector in $\mathbb{R}_+^\Omega$, 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 $\mathcal{A} : \mathcal{V} \to \Delta(\Omega)$ 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, $\frac{v^T \mathcal{A}(v)}{||v||_\infty} \geq c$. The following is an approximate variant of *weak monotonicity* [29], and follows from the fact that $\mathcal{A}$ is the allocation rule of a $(1-\epsilon)$-approximately truthful mechanism.

Fact A.1 (Similar to [29])**.** *For any $u, v \in \mathcal{V}$,*

$$v^T \mathcal{A}(v) - (1-\epsilon) u^T \mathcal{A}(v) \geq (1-\epsilon) v^T \mathcal{A}(u) - u^T \mathcal{A}(u).$$

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 $\mathcal{A}$.

Remark A.2. *We note that the proofs of this section apply more generally than CPP with submodular valuations. In particular, the only properties of this problem that are used in the proofs are: (1) The multiple-player allocation problem is algorithmically equivalent to the single player allocation problem (2) The set Ω of outcomes is finite, (3) The set of valuations $\mathcal{V} \subseteq \mathbb{R}_+^\Omega$ is closed under scaling by a non-negative constant, and (4) There is a deterministic algorithm $s : \mathcal{V} \to \mathbb{R}_+$ that runs in finite time, makes a polynomial number of value queries, and returns a "weak approximation" to the optimal value – i.e. we only require that $s(v) > 0$ when $||v||_\infty > 0$. When a welfare-maximization mechanism design problem satisfies these four conditions, as do all variants of CPP and other "public-project"-type problems in the literature, then the analogue of Theorem 4.1 holds for that problem.*

A.1 Limit behavior of truthful in expectation mechanisms

We will show that $\mathcal{A}$ is $(1-\epsilon)$-approximately MIDR in the limit as we scale up the valuations. Recall that a mechanism $\mathcal{B} : \mathcal{V} \to \Delta(\Omega)$ is $(1-\epsilon)$-approximately MIDR if $v^T \mathcal{B}(v) \geq (1-\epsilon) \sup_{w \in \mathcal{V}} v^T \mathcal{B}(w)$ for all $v \in \mathcal{V}$. The following statement is analogous.

Proposition A.3. $\liminf_{\alpha \to \infty} v^T \mathcal{A}(\alpha v) \geq (1-\epsilon) \sup_{w \in \mathcal{V}} v^T \mathcal{A}(w)$ *for all $v \in \mathcal{V}$.*

Proof. Let $\alpha, \beta \in \mathbb{R}_+$, and let $w, v \in \mathcal{V}$. By Fact A.1:

$$\alpha v^T \mathcal{A}(\alpha v) - (1-\epsilon) w^T \mathcal{A}(\alpha v) \geq (1-\epsilon) \alpha v^T \mathcal{A}(w) - w^T \mathcal{A}(w).$$

Dividing the expression by α, we get:

$$v^T \mathcal{A}(\alpha v) - \frac{(1-\epsilon) w^T \mathcal{A}(\alpha v)}{\alpha} \geq (1-\epsilon) v^T \mathcal{A}(w) - \frac{w^T \mathcal{A}(w)}{\alpha}$$

Taking the limit infimum as α goes to infinity,

$$\liminf_{\alpha\to\infty} v^T\mathcal{A}(\alpha v) \geq (1-\epsilon)v^T\mathcal{A}(w).$$

Now, taking the supremum over w

$$\liminf_{\alpha\to\infty} v^T\mathcal{A}(\alpha v) \geq (1-\epsilon)\sup_{w\in\mathcal{V}} v^T\mathcal{A}(w)$$

This completes the proof. □

A.2 Approximating the limit behavior of a mechanism

Ideally, we would transform $\mathcal{A}$ to an allocation rule that behaves as $\mathcal{A}$ does in the limit – by the results of the previous sub-section, 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.

For each $\delta' > 0$ and $v \in \mathcal{V}$, we define a threshold $t(\delta', v)$. Roughly speaking, $t(\delta', v)$ is the "scale" at which $\mathcal{A}$ is guaranteed to be within $(1-\delta')$ of its limit behavior when given input in the direction of v. Proposition A.3 guarantees that threshold $t(\delta', v)$ exists for each $v \in \mathcal{V}$ and $\delta' > 0$.

$$t(\delta', v) = \sup\left\{t : v^T\mathcal{A}\left(t\frac{v}{||v||_\infty}\right) \leq (1-\epsilon-\delta')\sup_{w\in\mathcal{V}} v^T\mathcal{A}(w)\right\} + 1 \tag{9}$$

We note that the motivation for adding 1 (any arbitrary positive number would do) to the expression is to guarantee that $v^T\mathcal{A}\left(t(\delta', v)\frac{v}{||v||_\infty}\right) \geq (1-\epsilon-\delta')\sup_{w\in\mathcal{V}} v^T\mathcal{A}(w)$, which may not be guaranteed by the supremum.

Assume that, for some $\delta' > 0$, we have some upper-bound τ on $\{t(\delta', v) : v \in \mathcal{V}\}$, and moreover we have a procedure $s(v)$ to estimate a non-zero lower bound on $||v||_\infty$ for each $v \in \mathcal{V}$. Then, the following procedure is evidently a c-approximate and $(1-\epsilon-\delta')$-MIDR allocation rule: Given input $v \in \mathcal{V}$, output $\mathcal{A}(\frac{\tau}{s(v)} \cdot v)$. The procedure $s(v)$ is easy to implement using only m value queries — indeed, we can take $s(v) = \max_{j\in[m]} v(\{j\})$. It is not clear, however, that the upper-bound τ can be computed effectively. Even worse, it is not clear that such an upper-bound even exists: $\mathcal{V}$ is infinite, and $t(\delta', v)$ is not necessarily a continuous function of v!

We remedy this as follows. We will show that there exists such an upper-bound when δ' is sufficiently large relative to ϵ. Specifically, we show that there exists an upper-bound τ on $\{t(\delta + 2\epsilon, v) : v \in \mathcal{V}\}$. However, computing such an upper-bound in finite time may be impossible in general, since the scale of valuations at which $\mathcal{A}$ approaches its limit behavior may be arbitrary. Instead, we take τ as advice to our non-uniform reduction. By the discussion in the previous paragraph, showing that the upper-bound τ exists yields a non-uniform allocation rule that is $(1-3\epsilon-\delta)$-MIDR, and makes at most m more value queries than $\mathcal{A}$, completing the proof of Theorem 4.1.

As a tool for proving that the upper-bound τ exists, we define a finite net of $\mathcal{V}$. Since $\mathcal{V}$ is a cone in finite-dimensional euclidean space, its intersection with the infinity-norm unit ball admits a *σ-net* in the infinity-norm for any $\sigma > 0$ — specifically, a finite set $\mathcal{U} \subseteq \mathcal{V}$ such that

1. $||u||_\infty = 1$ for all $u \in \mathcal{U}$
2. $\forall v \in \mathcal{V}\ \exists u \in \mathcal{U}\ \left|\left|\frac{v}{||v||_\infty} - u\right|\right|_\infty \leq \sigma$

Let $\sigma = c\delta/4$ and let $\mathcal{U}$ be a σ-net of $\mathcal{V}$. Now let $\beta = \max_{u\in\mathcal{U}} t(\frac{\delta}{4}, u)$, and let $\tau = 4\beta/\delta$. It suffices to show that $v^T\mathcal{A}(\tau v) \geq (1-3\epsilon-\delta)\sup_{w\in\mathcal{V}} v^T\mathcal{A}(w)$ for each $v \in \mathcal{V}$ with $||v||_\infty = 1$. Let $v \in \mathcal{V}$ be such that $||v||_\infty = 1$, and let u be a point in the σ-net $\mathcal{U}$ such that $||v-u||_\infty \leq \sigma$. By Fact A.1, we have

$$\tau v^T\mathcal{A}(\tau v) - (1-\epsilon)\beta u^T\mathcal{A}(\tau v) \geq (1-\epsilon)\tau v^T\mathcal{A}(\beta u) - \beta u^T\mathcal{A}(\beta u) \tag{10}$$

We now use inequality (10) to lower-bound $v^T\mathcal{A}(\tau v)$:

$$\begin{aligned}
v^T\mathcal{A}(\tau v) &\geq (1-\epsilon)v^T\mathcal{A}(\beta u) - \frac{\beta}{\tau}u^T\mathcal{A}(\beta u) && \text{Dividing (10) by } \tau \text{ and loosening the inequality}\\
&= (1-\epsilon)v^T\mathcal{A}(\beta u) - \frac{\delta}{4}u^T\mathcal{A}(\beta u) && \text{By definition of } \tau\\
&= \left(1-\epsilon-\frac{\delta}{4}\right)u^T\mathcal{A}(\beta u) - (1-\epsilon)(u-v)^T\mathcal{A}(\beta u)\\
&\geq \left(1-\epsilon-\frac{\delta}{4}\right)u^T\mathcal{A}(\beta u) - ||u-v||_\infty && \text{Since } ||\mathcal{A}(\beta u)||_1 = 1\\
&\geq \left(1-\epsilon-\frac{\delta}{4}\right)u^T\mathcal{A}(\beta u) - \sigma && \text{By proximity of } u \text{ and } v\\
&\geq \left(1-\epsilon-\frac{\delta}{2}\right)u^T\mathcal{A}(\beta u) && \text{By } c \leq u^T\mathcal{A}(\beta u) \text{ and definition of } \sigma\\
&\geq \left(1-2\epsilon-\frac{3}{4}\delta\right)\sup_{w\in\mathcal{V}} u^T\mathcal{A}(w) && \text{By definition of } \beta\\
&\geq \left(1-2\epsilon-\frac{3}{4}\delta\right)\liminf_{\alpha\to\infty} u^T\mathcal{A}(\alpha v)\\
&= \left(1-2\epsilon-\frac{3}{4}\delta\right)\liminf_{\alpha\to\infty}(v^T\mathcal{A}(\alpha v) - (v-u)^T\mathcal{A}(\alpha v))\\
&\geq \left(1-2\epsilon-\frac{3}{4}\delta\right)\liminf_{\alpha\to\infty}(v^T\mathcal{A}(\alpha v) - ||v-u||_\infty) && \text{Since } ||\mathcal{A}(\alpha v)||_1 = 1\\
&\geq \left(1-2\epsilon-\frac{3}{4}\delta\right)\liminf_{\alpha\to\infty}(v^T\mathcal{A}(\alpha v) - \sigma) && \text{By proximity of } u \text{ and } v\\
&\geq \left(1-2\epsilon-\frac{3}{4}\delta\right)\liminf_{\alpha\to\infty}(v^T\mathcal{A}(\alpha v) - \frac{\delta}{4}v^T\mathcal{A}(\alpha v)) && \text{By } c \leq v^T\mathcal{A}(\alpha v) \text{ and definition of } \sigma\\
&\geq (1-2\epsilon-\delta)\liminf_{\alpha\to\infty} v^T\mathcal{A}(\alpha v)\\
&\geq (1-3\epsilon-\delta)\sup_{w\in\mathcal{V}} v^T\mathcal{A}(w) && \text{By Proposition A.3}
\end{aligned}$$

By the previous discussion, this completes the proof of Theorem 4.1.

B From approximately-MIDR to approximately-TIE mechanisms for combinatorial auctions

In this section, we show how an approximately-MIDR allocation rule can be used to design an approximately-TIE mechanism for combinatorial auctions. In particular, we show how the mechanism of Dughmi, Roughgarden and Yan [20] for combinatorial auctions can be implemented in the value oracle model, at the cost of replacing the property of truthfulness in expectation by approximate truthfulness in expectation.

The mechanism of [20] is presented in a "lottery-value" oracle model, where each bidder can be queried about his valuation by means of the following query: given a vector of probabilities over items $\mathbf{x} \in [0,1]^m$,

what is the expected value $\mathbf{E}[v_i(\widehat{\mathbf{x}})]$, where $\widehat{\mathbf{x}}$ is obtained by independently rounding each coordinate of $\mathbf{x}$ to 0 or 1 with probability x_i. Such queries can be answered efficiently for certain valuation functions (in particular *coverage* functions), and this oracle model is a convenient framework for the presentation of the mechanism of [20]. On the other hand, lottery-value queries are #P-hard to answer for the class of matroid rank functions (see Section C), and hence one can ask how realistic this model is in general. The purpose of this section is to show that the model is "approximately realistic" in the sense that the mechanism of [20] can be implemented in the (weaker) value oracle model at the cost of relaxing the concept of truthfulness in expectation to *approximate truthfulness in expectation*. (Here, we mean appoximation within an arbitrarily small error, in the sense of an FPTAS.) In particular, we show that the maximal in distributional range (MIDR) allocation rule of [20] can be implemented as an approximately MIDR allocation rule in the value oracle model, and then we present a blackbox transformation from approximately MIDR allocation rules to approximately TIE mechanisms.

The class of valuations of interest here is the following (as in [20]).

Definition B.1. *A function $v : 2^{[m]} \to \mathbb{R}_+$ is a* weighted matroid rank sum *if there are matroids $\mathcal{M}_1, \ldots, \mathcal{M}_k$ and weights $\alpha_1, \ldots, \alpha_k \geq 0$ such that*

$$v(S) = \sum_{i=1}^{k} \alpha_i r_{\mathcal{M}_i}(S),$$

where $r_{\mathcal{M}_i}$ is the rank function of matroid $\mathcal{M}_i$.

This definition also captures positive combinations of *weighted rank functions*, as the cones generated by weighted and unweighted rank functions of matroids coincide.

The allocation rule of [20] is maximal-in-distributional-range, provides a $(1 - 1/e)$-approximation to the social welfare, and the corresponding TIE mechanism can be implemented in expected polynomial time provided the bidders' valuations are weighted matroid rank sums and support lottery-value queries. Our goal here is to prove the following.

Theorem B.2. *For every $\epsilon = 1/poly(m, n)$, there is a $(1-\epsilon)$-TIE mechanism that achieves a $(1-1/e-\epsilon)$-approximation to the social welfare in combinatorial auctions with weighted-matroid-rank-sum valuations that runs in polynomial time in the value oracle model.*

We prove this claim in two steps. First, we prove the following.

Theorem B.3. *For every $\epsilon = 1/poly(m, n)$, there is a $(1-\epsilon)$-MIDR allocation rule that achieves a $(1-1/e-\epsilon)$-approximation to the social welfare in combinatorial auctions with weighted-matroid-rank-sum bidders and runs in polynomial time in the value oracle model.*

Then, we prove the following general reduction.

Theorem B.4. *For every $(1-\epsilon)$-MIDR allocation rule that is a c-approximation to the social welfare in combinatorial auctions with bidders' valuations restricted to a set $\mathcal{C}$, there is a $(1-\epsilon')$-TIE mechanism that is a $(c - 1/poly(m, n))$-approximation to the social welfare in combinatorial auctions with bidders' valuations in $\mathcal{C}$, where $\epsilon' = \epsilon \cdot poly(m, n)$.*

Therefore, by selecting a sufficiently (polynomially) small ϵ in Theorem B.3, we can achieve an arbitrarily (polynomially) small error ϵ in Theorem B.2.

B.1 An approximately MIDR allocation rule

In this section, we prove Theorem B.3. We use a variant of the mechanism of [20]. Instead of sophisticated convex optimization techniques, which seem necessary to find the exact optimum over the range, we use a simple local search that guarantees that we get arbitrarily close to the optimum. We begin with some definitions.

Definition B.5. *For a combinatorial auction with m items and n bidders with valuations $v_i : 2^{[m]} \to \mathbb{R}_+$, the* aggregate valuation function $f : 2^{[n]\times[m]} \to \mathbb{R}_+$ *is*

$$f(S) = \sum_{i=1}^{n} v_i(\pi_i(S)),$$

where $\pi_i(S) = \{j : (i,j) \in S\}$. We define $F : [0,1]^{[n]\times[m]} \to \mathbb{R}_+$ to be the multilinear extension of f (see also [39]), and P to be the polytope of fractional allocations:

$$P = \left\{ \mathbf{x} \in [0,1]^{[n]\times[m]} : \forall j : \sum_{i=1}^{n} x_{ij} \leq 1 \right\}.$$

The (integral) welfare maximization problem turns out to be equivalent to $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$. This problem cannot be solved optimally, even for very special classes of valuation functions. In lieu of F, the authors of [20] use the modified objective function

$$F^{exp}(x_{11}, x_{12}, \ldots, x_{nm}) = F(1 - e^{-x_{11}}, 1 - e^{-x_{12}}, \ldots, e^{-x_{nm}}).$$

Interestingly, the function F^{exp} turns out to be *concave* for a subclass of submodular functions, including weighted matroid rank sums (see [20] for a proof). This means that we can solve the problem $\max\{F^{exp}(\mathbf{x}) : \mathbf{x} \in P\}$, which means in effect optimizing over a certain range of product distributions. Also, the optimum of this problem is at least $(1 - 1/e)$ times the optimal social welfare. Supplementing this MIDR allocation rule with suitable payments yields a $(1 - 1/e)$-approximate, TIE mechanism [20].

Here, we propose the following simple algorithm that solves the problem $\max\{F^{exp}(\mathbf{x}) : \mathbf{x} \in P\}$ near-optimally (in the sense of an FPTAS).

Local Search Allocation Rule.

- Initialize $\mathbf{x} := 0$. Let M be the maximum value of any singleton.
- Let $\mathbf{g}$ be an estimate of the gradient $\nabla F^{exp}(\mathbf{x})$, within additive error δM in each coordinate, where $\delta = \frac{\epsilon}{8m^2n^2}$. As long as there is a point $\mathbf{y} \in P$ such that

  $$(\mathbf{y} - \mathbf{x}) \cdot \mathbf{g} > \frac{1}{2}\epsilon M,$$

  update $\mathbf{x} := \mathbf{x} + \delta(\mathbf{y} - \mathbf{x})$.
- Return an allocation randomly sampled from the distribution $\mathbf{x}$.

The required estimates of $\nabla F^{\exp}(\mathbf{x})$ can be obtained in polynomial time by random sampling, with high probability. (By *high probability*, we mean $1 - e^{-poly(m,n)}$ in this section. The coordinates of $\nabla F^{exp}(\mathbf{x})$ are always in the interval $[0, M]$ — by submodularity — and so this follows from standard Chernoff bounds [1].) Linear programming can be used to efficiently find a suitable point $\mathbf{y}$, or certify that no such point exists.

B.2 The analysis

We claim that this allocation rule runs in polynomial time and solves the problem $\max\{F^{exp}(\mathbf{x}) : \mathbf{x} \in P\}$ up to a $(1 - \epsilon)$ factor with high probability, thus proving Theorem B.3. In the following, we assume that the estimate of $\nabla F^{exp}(\mathbf{x})$ obtained in each step is accurate within additive error δM, which is possible to achieve with high probability over the run of the algorithm, via a polynomial number of samples.

We proceed via a series of claims.

Lemma B.6. *If the algorithm terminates, then with high probability*

$$F^{exp}(\mathbf{x}) \geq (1 - \epsilon) \max\{F^{exp}(\mathbf{x}) : \mathbf{x} \in P\}.$$

Proof. Let $\mathbf{y}$ be an optimal solution of $\max\{F^{exp}(\mathbf{x}) : \mathbf{x} \in P\}$. When the algorithm terminates at $\mathbf{x}$, we have $(\mathbf{y}-\mathbf{x}) \cdot \nabla F^{exp}(\mathbf{x}) < \epsilon M$ (even accounting for the errors in our estimate of $\nabla F^{exp}(\mathbf{x})$). By the concavity of F^{exp},

$$OPT - F^{exp}(\mathbf{x}) = F^{exp}(\mathbf{y}) - F^{exp}(\mathbf{x}) \leq (\mathbf{y}-\mathbf{x}) \cdot \nabla F^{exp}(\mathbf{x}) \leq \epsilon M \leq \epsilon OPT.$$

□

Lemma B.7. *In each iteration, with high probability, the value of $F^{exp}(\mathbf{x})$ increases by at least $\frac{\epsilon^2}{64m^2n^2}M$.*

Proof. If the algorithm continues, we can assume that $(\mathbf{y}-\mathbf{x}) \cdot \nabla F^{exp}(\mathbf{x}) > \frac{1}{4}\epsilon M$ (considering that the estimate of $\nabla F^{exp}(\mathbf{x})$ could be off by $\delta M = \frac{\epsilon M}{8m^2n^2}$ in each coordinate). We also have bounds on how much the gradient can change when $\mathbf{x}$ moves by a certain amount. Specifically, for $(i,j) \neq (i',j')$,

$$\Big|\frac{\partial^2 F^{exp}}{\partial x_{ij}\partial x_{i'j'}}\Big| = \Big|e^{-x_{ij}-x_{i'j'}}\frac{\partial^2 F}{\partial x_{ij}\partial x_{i'j'}}\Big| \leq \Big|\frac{\partial^2 F}{\partial x_{ij}\partial x_{i'j'}}\Big| \leq M$$

and similarly

$$\Big|\frac{\partial^2 F^{exp}}{\partial {x_{ij}}^2}\Big| = \Big|e^{-x_{ij}}\frac{\partial F}{\partial x_{ij}}\Big| \leq \Big|\frac{\partial F}{\partial x_{ij}}\Big| \leq M,$$

by known properties of the multilinear extension [39, 40]. This implies that for any $\mathbf{x}'$ such that $||\mathbf{x}'-\mathbf{x}||_\infty \leq \delta$,

$$\frac{\partial F^{exp}}{\partial x_{ij}}\Big|_{\mathbf{x}'} \geq \frac{\partial F^{exp}}{\partial x_{ij}}\Big|_{\mathbf{x}} - \sum_{i,j}|x'_{ij}-x_{ij}| \max\Big|\frac{\partial^2 F^{exp}}{\partial x_{ij}\partial x_{i'j'}}\Big| \geq \frac{\partial F^{exp}}{\partial x_{ij}}\Big|_{\mathbf{x}} - \delta mnM.$$

Consequently,

$$\begin{aligned} F^{exp}(\mathbf{x}+\delta(\mathbf{y}-\mathbf{x})) &\geq F^{exp}(\mathbf{x}) + \delta(\mathbf{y}-\mathbf{x}) \cdot (\nabla F^{exp}(\mathbf{x}) - \delta mnM\mathbf{1}) \\ &\geq F^{exp}(\mathbf{x}) + \delta(\mathbf{y}-\mathbf{x}) \cdot \nabla F^{exp}(\mathbf{x}) - \delta^2 m^2 n^2 M \\ &\geq F^{exp}(\mathbf{x}) + \delta \cdot \frac{1}{4}\epsilon M - \delta^2 m^2 n^2 M. \end{aligned}$$

Again using $\delta = \frac{\epsilon}{8m^2n^2}$, we obtain

$$F^{exp}(\mathbf{x}+\delta(\mathbf{y}-\mathbf{x})) \geq F^{exp}(\mathbf{x}) + \frac{\epsilon^2}{32m^2n^2}M - \frac{\epsilon^2}{64m^2n^2}M \geq F^{exp}(\mathbf{x}) + \frac{\epsilon^2}{64m^2n^2}M.$$

□

Lemma B.8. *The number of iterations is with high probability at most $64m^3n^2/\epsilon^2$.*

Proof. By the previous lemma, with high probability the value of $F^{exp}(\mathbf{x})$ increases in each iteration by at least $\frac{\epsilon^2}{64m^2n^2}M$. After $64m^3n^2/\epsilon^2$ iterations, it will be at least mM. By the definition of M and submodularity of valuations, mM is an upper bound on the welfare of every feasible allocation, and hence also of the function F^{exp}. This completes the proof. □

This concludes the proof of Theorem B.3. We remark that, building on the allocation rule in [22], a similar approach gives a $(1-\epsilon)$-MIDR and $(1-1/e-\epsilon)$-approximate allocation rule for combinatorial public projects with weighted matroid rank sums that runs in polynomial time in the value oracle model.

B.3 From approximately MIDR to approximately TIE

In this section, we prove Theorem B.4. We assume that we have a $(1-\epsilon)$-MIDR allocation rule $\mathcal{M}$ providing a c-approximation for combinatorial auctions with valuations in a class $\mathcal{C}$. We assume in the following that $c \geq \frac{1}{n}$, where n is the number of bidders. (A $\frac{1}{n}$-approximation is trivial to achieve by giving all of the items to a random bidder.) We also assume that $\epsilon = 1/poly(m,n)$. We want to convert the $(1-\epsilon)$-MIDR allocation rule into an $(1-\epsilon')$-TIE mechanism. Our approach is as follows. If $\epsilon = 0$, then the VCG payment scheme turns an MIDR mechanism into a TIE mechanism. The fact that our mechanism is only approximately MIDR means that the VCG payments might suffer from errors that are significant for certain bidders, especially if their utility is close to zero. Therefore, we modify the mechanism to ensure that bidders whose valuation is very low do not participate in the VCG scheme. In addition, we provide each bidder with the bundle of *all* items with some small probability, so that their expected utility is not extremely small. Our mechanism works as follows.

Mechanism $\mathcal{M}'$.

1. Let $\mathcal{M}$ be a $(1-\epsilon)$-MIDR allocation rule.
2. Let V_i be the valuation that bidder i reports for the ground set of all items. Run $\mathcal{M}$ to compute a distribution over allocations $(S_1, \ldots, S_n)$ and let O_i be an (unbiased) estimate of the expected value collected by bidder i, $\mathbf{E}[v_i(S_i)]$.[20] Let $O = \sum_{i=1}^n O_i$ be an estimate of the expected social welfare $\sum_{i=1}^n \mathbf{E}[v_i(S_i)]$. For each bidder i, run $\mathcal{M}$ also on the same instance without bidder i, and denote by O'_{-i} an estimate of the expected social welfare of its outcome.
3. Call bidder i *relevant* if
$$V_i > \frac{1}{n^7} \sum_{j \neq i} V_j.$$
Call bidder i *active* if he is relevant and, in addition,
$$\left(1 - \frac{1}{n}\right)(O - O'_{-i}) + \frac{1}{2n^2} V_i > \frac{1}{n^4} O'_{-i}.$$
4. With probability $1 - 1/n$: allocate a set S_i from the distribution found by $\mathcal{M}$ to each active bidder i, and charge the VCG-like price $p_i = O'_{-i} - \sum_{j \neq i} O_j$. Do not allocate or charge anything to inactive bidders.
5. Else, with probability $1/n^2$ for each bidder i: If active, allocate the ground set to i and charge $\frac{1}{n^2} O'_{-i}$. If inactive, allocate the ground set with probability $1/2$ and charge 0.

We emphasize that O and O'_{-i} are random variables that we obtain by running the (randomized) allocation rule $\mathcal{M}$. We denote the actual optima over the range, with respect to the reported valuations, by OPT and OPT'_{-i}. In expectation, we have $OPT \geq \mathbf{E}[O] \geq (1-\epsilon)OPT$ and $OPT'_{-i} \geq \mathbf{E}[O'_{-i}] \geq (1-\epsilon)OPT'_{-i}$; however with some probability, O could be significantly different from OPT (even larger, since it is a probabilistic estimate), and O'_{-i} could be significantly different from OPT'_{-i}.

In the following, we denote by v_i^* the actual valuation of bidder i, and by $V_i^* = v_i^*(M)$ the actual value of the ground set for bidder i. Let OPT_{+i} denote the optimum over the range with valuations v_j for $j \neq i$ and v_i^* for bidder i. (Note that $OPT_{+i} = OPT$ if $v_i^* = v_i$.) Let O_{+i} denote our estimate of OPT_{+i} (assuming bidder i reports the truth). We prove the following statements.

Lemma B.9. *For every bidder such that $V_i^* \leq \frac{1}{n^7} \sum_{j \neq i} V_j$, his expected utility is maximized within an ϵ-fraction of his utility by reporting truthfully.*

[20] By polynomially bounded sampling, we can assume that our estimate O_i is with high probability within $\mathbf{E}[v_i(S_i)] \pm V_i/poly(m,n)$, and the probability of deviation decays exponentially. Similarly for the estimates of O and O'_{-i}.

Proof. Observe that as long as bidder i reports $V_i \leq \frac{1}{n^7}\sum_{j\neq i} V_j$, he is inactive and receives the same utility regardless of his bid. Therefore, his utility could change only if he reports $V_i > \frac{1}{n^7}\sum_{j\neq i} V_j$. In that case, he might be classified as active (depending on O and O'_{-i}). However, if that happens then he is charged at least $\frac{1}{n^2}O'_{-i}$ with probability $\frac{1}{n^2}$, i.e. $\frac{1}{n^4}O'_{-i}$ in expectation (conditioned on the value of O'_i). Since the most value he can ever collect is V_i^*, he would (possibly) gain from being active only if $O'_{-i} < n^4 V_i^* \leq \frac{1}{n^3}\sum_{j\neq i} V_j$. Since $\mathbf{E}[O'_{-i}] \geq (1-\epsilon)OPT'_{-i} > \frac{1}{n^2}\sum_{j\neq i} V_j$, it is very unlikely that O'_{-i} is less than $\frac{1}{n^3}\sum_{j\neq i} V_j$; this happens with exponentially small probability. Hence, by lying, bidder i could possibly gain only an exponentially small fraction of V_i^* in expectation, negligible with respect to his utility as an inactive player. □

Lemma B.10. *For every bidder i such that $V_i^* > \frac{1}{n^7}\sum_{j\neq i} V_j$, we have*

$$\mathbf{E}[|OPT_{+i} - O_{+i}|] \leq 3\epsilon n^7 V_i^*$$

and

$$\mathbf{E}[|OPT'_{-i} - O'_{-i}|] \leq 2\epsilon n^7 V_i^*.$$

Proof. We have $V_i^* > \frac{1}{n^7}\sum_{j\neq i} V_j \geq \frac{1}{n^7}OPT'_{-i}$. We also know that $|OPT'_{-i} - \mathbf{E}[O'_{-i}]| \leq \epsilon OPT'_{-i} \leq \epsilon n^7 V_i^*$. The estimate O'_{-i} of the output of the mechanism for all bidders except i is concentrated around its expectation $\mathbf{E}[O'_{-i}]$, with variance $\frac{1}{poly(m,n)}\sum_{j\neq i} V_j \ll \epsilon n^7 V_i^*$, hence we can estimate

$$\mathbf{E}[|OPT'_{-i} - O'_{-i}|] \leq 2\epsilon n^7 V_i^*.$$

Similarly, $|OPT_{+i} - \mathbf{E}[O_{+i}]| \leq \epsilon OPT_{+i} \leq \epsilon(OPT_{-i} + V_i^*) \leq 2\epsilon n^7 V_i^*$, and O_{+i} is concentrated around its expectation, therefore $\mathbf{E}[|OPT_{+i} - O_{+i}|] \leq 3\epsilon n^7 V_i^*$. □

Lemma B.11. *Every bidder i such that $V_i^* > \frac{1}{n^7}\sum_{j\neq i} V_j$ maximizes his expected utility within a factor of $(1 - O(\epsilon n^9))$ by reporting his true valuation.*

Proof. Let us fix the valuations of all bidders except i. Let us assume for now that our estimate O_j is exactly equal to $\mathbf{E}[v_j(S_j)]$, $O = \sum_{j=1}^n O_j$ is equal to OPT (meaning that the MIDR mechanims optimizes exactly), and similarly O'_{-i} is equal to OPT'_{-i}. We will analyze this idealized mechanism first.

If bidder i ends up being active, his expected utility will be

$$\begin{aligned} U_{active} &= (1-1/n)(\mathbf{E}[v_i^*(S_i)] - (O'_{-i} - \sum_{j\neq i} O_j)) + \frac{1}{n^2}V_i^* - \frac{1}{n^4}OPT'_{-i} \\ &= (1-1/n)(\mathbf{E}[v_i^*(S_i)] + \sum_{j\neq i}\mathbf{E}[v_j(S_j)] - OPT'_i) + \frac{1}{n^2}V_i^* - \frac{1}{n^4}OPT'_{-i} \\ &\leq (1-1/n)(OPT_{+i} - OPT'_{-i}) + \frac{1}{n^2}V_i^* - \frac{1}{n^4}OPT'_{-i} = U^+_{active}. \end{aligned}$$

Here we used the fact that OPT_{+i} is the optimal value over the range with valuations v_i^* and $v_j, j \neq i$. This implies that the last quantity, U^+_{active}, is the best possible utility bidder i could receive as an active bidder. In fact he will receive this utility if he reports truthfully and ends up being active.

If bidder i is inactive, then his expected utility will be

$$U_{inactive} = \frac{1}{2n^2}V_i^*.$$

Now, if it is the case that

$$U^+_{active} - U_{inactive} = (1-1/n)(OPT_{+i} - OPT'_{-i}) + \frac{1}{2n^2}V_i^* - \frac{1}{n^4}OPT'_{-i} \leq 0,$$

this means that no matter what bidder i reports, being active cannot be more profitable than not being active for him. When reporting his true valuation, such a bidder will in fact be inactive, because the condition for

making a bidder active is exactly $(1-1/n)(O-O'_{-i})+\frac{1}{2n^2}V_i-\frac{1}{n^4}O'_{-i}>0$, and in this case we would have $OPT_{+i}=O$ and $O'_{-i}=OPT'_{-i}$. Other than making the bidder inactive, the particular valuation he reports doesn't have an impact on his utility, so he might as well report the truth.

On the other hand, if

$$U^+_{active}-U_{inactive}=(1-1/n)(OPT_{+i}-OPT'_{-i})+\frac{1}{2n^2}V_i^*-\frac{1}{n^4}OPT'_{-i}>0,$$

then it is more profitable for bidder i to be active, since by reporting truthfully he will get utility U^+_{active}, better than $U_{inactive}$ as an inactive bidder. In fact we argued above that an active bidder cannot get a better utility by reporting any valuation, so the best strategy for him is to report truthfully. In conclusion, the idealized mechanism rewards a truthfully reporting bidder by utility $\max\{U^+_{active},U_{inactive}\}$ and that's the best the bidder can possibly receive.

Finally, we have to deal with the fact that $O=\sum_{j=1}^n O_j$ is not exactly equal to OPT, and O'_{-i} is not exactly equal to OPT'_{-i}. By Lemma B.10, the estimates O_{+i} and O'_{-i} are in expectation within $O(\epsilon n^7 V_i^*)$ of OPT_{+i} and OPT_{-i}. The estimates $\sum_{j\neq i}O_j$ of $\sum_{j\neq i}\mathbf{E}[v_j(S_j)]$ are strongly concentrated, let's say with high probability within $\epsilon\sum_{j\neq i}V_j=O(\epsilon n^7V_i^*)$ of the expectation. Also, the actual social welfare of the distribution returned by the mechanism when bidder i reports truthfully is within $O(\epsilon n^7V_i^*)$ of OPT_{+i} in expectation. Therefore, the expected utility of a truthfully reporting bidder is at least $\max\{U^+_{active},U_{inactive}\}-O(\epsilon n^7V_i^*)$. On the other hand, the expected utility under any other reported valuation cannot be better than $\max\{U^+_{active},U_{inactive}\}+O(\epsilon n^7V_i^*)$, again due to the precision of the estimates stated above. We have also ensured that the bidder's utility is at least $\frac{1}{2n^2}V_i^*$. Therefore, the relative error in utility maximization is $O(\epsilon n^9)$. □

Now we can prove Theorem B.4.

Proof. Using a $(1-\epsilon)$-MIDR mechanism $\mathcal{M}$, we implement a new mechanism $\mathcal{M}'$ as above. By Lemma B.9 and B.11, each bidder maximizes his utility within a factor of $1-O(\epsilon\cdot poly(n))$ by reporting truthfully. Moreover, the expected social welfare provided by mechanism $\mathcal{M}'$ is at least $(1-1/n)$ times the social welfare of all *active* bidders in $\mathcal{M}$. (Just considering the option that we used the VCG-based allocation.) It remains to estimate the loss in social welfare due to inactive bidders.

Consider a bidder i such that $V_i\geq\frac{4}{n^2}OPT'_{-i}$. Since $OPT'_{-i}\geq\mathbf{E}[O'_{-i}]$, and O'_{-i} is a strongly concentrated estimate, with high probability we also have $V_i\geq\frac{2}{n^2}O'_{-i}$. Therefore, with high probability the bidder will be active and participate in the VCG scheme. The only bidders who do not participate with significant probability are those such that $V_i<\frac{4}{n^2}OPT'_{-i}\leq\frac{4}{n^2}OPT$. Therefore, all such bidders together cannot amount to more than $\frac{4}{n}OPT$. Overall, we recover at least a $(1-O(1/n))$-fraction of the social welfare achieved by mechanism $\mathcal{M}$, which means an approximation factor at least $(1-O(1/n))c$. It is easy to see that the $O(1/n)$ term can be replaced by any inverse polynomial in m,n, if desired. □

C #P-hardness of lottery value queries

Here we show that for matroid rank functions, lottery-value queries are #P-hard to answer, and require an exponential number of queries if the matroid is given by an independence oracle. We note that a lottery-value query for the vector $\mathbf{x}=(\frac{1}{2},\ldots,\frac{1}{2})$ is simply an expectation over a uniformly random set of elements.

Theorem C.1. *There is a class of succintly represented matroids for which it is #P-hard to compute* $\mathbf{E}[r_{\mathcal{M}}(R)]$*, where* $r_{\mathcal{M}}$ *is the rank function of* $\mathcal{M}$ *and* R *is a uniformly random set of matroid elements. For matroids given by an independence oracle, computing* $\mathbf{E}[r_{\mathcal{M}}(R)]$ *requires an exponential number of queries.*

Proof. We use the class of "paving matroids" [37]: Let E be a ground set of size $2m$ partitioned into disjoint pairs $e_1,e_2,\ldots,e_m$, and let $\mathcal{F}\subset\binom{[m]}{k}$ be any family of k-element subsets of $[m]$. Then the following is a matroid: $S\subseteq E$ is independent iff either $|S|<2k$, or $|S|=2k$ and S is *not* a union of pairs $\bigcup_{i\in F}e_i$ where $F\in\mathcal{F}$.

Using this construction, we can embed any #P-hard problem in a paving matroid $\mathcal{M}$. For example, consider the problem of counting perfect matchings. For a graph G with m edges and n vertices, we let $k = n/2$ and we define $\mathcal{F}$ to be the family of k-edge subsets of edges that form a perfect matching. Then the matroid $\mathcal{M}$ defined as above captures the structure of perfect matchings, since for any set of edges F the rank function $r_{\mathcal{M}}(F)$ tells us whether F is a perfect matching (which is the case if and only if $r_{\mathcal{M}}(\bigcup_{i\in F} e_i) = 2|F| - 1 = 2k - 1$). Also, the matroid is succintly represented by the graph G, in the sense that given G we can easily decide whether a given set is independent in $\mathcal{M}$ or not. The value of $r_{\mathcal{M}}(S)$ depends on the structure of S only if S is a union of k pairs e_i, otherwise it is $r_{\mathcal{M}}(S) = \min\{|S|, 2k\}$. Therefore, if we can compute the value of $\mathbf{E}[r_{\mathcal{M}}(R)] = 2^{-2m}\sum_{S\subseteq[2m]} r_{\mathcal{M}}(S)$, we can extract the number of perfect matchings by an elementary formula.

Similarly, for a paving matroid given by an independence oracle, the value of $\mathbf{E}[r_{\mathcal{M}}(R)]$ determines the size of the family $\mathcal{F}$, which could be any arbitrary family. We cannot compute this value unless we determine the size of $\mathcal{F}$, which requires querying all sets of the form $\bigcup_{i\in F} e_i$. This requires an exponential number of queries for independence in $\mathcal{M}$. □

D Chernoff bound for bisections

Lemma D.1. *Suppose S is a fixed subset of $[m']$, and (A, B) a random partition of $[m']$, chosen uniformly among all partitions where $|A| = |B| = m'/2$. Then*

$$\Pr\left[||S \cap A| - |S \cap B|| > \beta m'\right] < 4e^{-\beta^2 m'/2}.$$

Proof. We use the fact that A has distribution very close to a uniformly random subset of $[m']$ (where elements appear independently with probability $1/2$). More precisely, we couple the two distributions as follows. Let A be a random set of size $m'/2$, B its complement, and let X be a binomial random variable $Bi(m', 1/2)$. Let R be a random set chosen as follows: if $X \le m'/2$, take a random subset of A of size X. If $X > m'/2$, take the union of A and $X - m'/2$ random elements from B. This defines a set R which is uniformly random. Hence, by the Chernoff bound (see e.g. [1, Theorem A.1.16]),

$$\Pr[|R\Delta A| > \alpha m'] = \Pr[Bi(m', 1/2) \notin [m'/2 - \alpha m', m'/2 + \alpha m']] < 2e^{-2\alpha^2 m'}.$$

Similarly, $S\Delta R$ has the distribution of a uniformly random set (because S is fixed), and hence

$$\Pr[|S\Delta R| \notin [m'/2 - \alpha m', m'/2 + \alpha m']] < 2e^{-2\alpha^2 m'}.$$

Using the triangle inequality $|S\Delta A| \le |S\Delta R| + |R\Delta A|$, we get

$$\Pr[|S\Delta A| \notin [m'/2 - 2\alpha m', m'/2 + 2\alpha m']] < 4e^{-2\alpha^2 m'}.$$

The lemma follows by taking $\alpha = \beta/2$, since $|S \cap A| - |S \cap B| = |A| - |S\Delta A| = \frac{m'}{2} - |S\Delta A|$. □

E Product composition of submodular functions

Lemma E.1. *Let $f_1, f_2 : 2^M \to [0, 1]$ be monotone submodular. Then*

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

is also monotone submodular.

Proof. Let $g_1(S) = 1 - f_1(S)$, $g_2(S) = 1 - f_2(S)$; these are non-negative non-increasing supermodular functions. Clearly, $g(S) = g_1(S)g_2(S)$ is also non-increasing. Our goal is to prove that $g(S) = g_1(S)g_2(S)$ is supermodular, which implies the claim. By the properties of g_1, g_2, we get for any $i, j \notin S$

$$g_1(S)(g_2(S) - g_2(S + i)) \ge g_1(S)(g_2(S + j) - g_2(S + i + j)) \ge g_1(S + j)(g_2(S + j) - g_2(S + i + j))$$

and

$$(g_1(S) - g_1(S+i))g_2(S+i) \geq (g_1(S+j) - g_1(S+i+j))g_2(S+i) \geq (g_1(S+j) - g_1(S+i+j))g_2(S+i+j).$$

Adding up these two inequalities, we get the condition of supermodularity for $g(S) = g_1(S)g_2(S)$:

$$g_1(S)g_2(S) - g_1(S+i)g_2(S+i) \geq g_1(S+j)g_2(S+j) - g_1(S+i+j)g_2(S+i+j).$$

□