

Fine-Grained Buy-Many Mechanisms Are Not Much Better Than Bundling

Sepehr Assadi^{*†}
sepehr.assadi@rutgers.edu

Vikram Kher[‡]
vkher@usc.edu

George Li[§]
gzli929@umd.edu

Ariel Schvartzman^{*¶}
aschvartzman@google.com

Abstract

Multi-item revenue-optimal mechanisms are known to be extremely complex, often offering buyers randomized lotteries of goods. In the standard buy-one model, it is known that optimal mechanisms can yield revenue infinitely higher than that of any “simple” mechanism—the ones with size polynomial in the number of items—even with just two items and a single buyer [4, 19].

We introduce a new parameterized class of mechanisms, buy- k mechanisms, which smoothly interpolate between the classical buy-one mechanisms and the recently studied buy-many mechanisms [12, 14, 13, 11]. Buy- k mechanisms allow the buyer to buy up to k many menu options. We show that restricting the seller to the class of buy- n incentive-compatible mechanisms suffices to overcome the bizarre, infinite revenue properties of the buy-one model. Our main result is that the revenue gap with respect to bundling, an extremely simple mechanism, is bounded by $O(n^2)$ for any arbitrarily correlated distribution $\mathcal{D}$ over n items for the case of an additive buyer. Our techniques also allow us to prove similar upper bounds for arbitrary monotone valuations, albeit with an exponential factor in the approximation.

On the negative side, we show that allowing the buyer to purchase a small number of menu options *does not suffice* to guarantee sub-exponential approximations, even when we weaken the benchmark to the optimal buy- k *deterministic* mechanism. If an additive buyer is only allowed to buy $k = \Theta(n^{1/2-\epsilon})$ many menu options, the gap between the revenue-optimal deterministic buy- k mechanism and bundling may be exponential in n . In particular, this implies that no “simple” mechanism can obtain a sub-exponential approximation in this regime.

1 Introduction

How should a revenue-maximizing seller price an item for sale when facing a buyer with a private value for the item? If the seller knows the distribution of values, seminal work of Myerson [24]

^{*}The authors would like to thank S. Matthew Weinberg and anonymous reviewers for providing feedback on the presentation of this paper.

[†]Department of Computer Science, Rutgers University. Research supported in part by a NSF CAREER Grant CCF-2047061, and a gift from Google Research.

[‡]University of Southern California. This work was carried out while Vikram Kher was a participant in the 2022 DIMACS REU program at Rutgers University, supported by NSF grants CCF-1852215 and CNS-2150186.

[§]University of Maryland. This work was carried out while George Li was a participant in the 2022 DIMACS REU program at Rutgers University, supported by NSF grants CCF-1852215 and CNS-2150186.

[¶]Google Research. Part of this research conducted while the author was at DIMACS, Rutgers University and was supported by the National Science Foundation under grant number CCF-1445755.

showed that it is optimal for the seller to offer the item at a take-it-or-leave-it price. The answer to this question becomes unclear for the case of multiple, even just two, items.

Optimal multi-item auctions are known to be complex objects, offering no discernible mathematical structure and often exhibiting “intuition-defying” properties [15]. A particularly egregious one is that there exist correlated distributions over just two items such that the revenue-optimal mechanism is *infinitely* better than any “simple” mechanism, ruling out the possibility of good worst-case approximations [18, 19].¹ The bizarre aspect of these pathological distributions is that the optimal revenue is *unbounded*, but any finite-sized mechanism can get at most finite revenue. One possible explanation for this bizarre phenomenon is that the seller is unrestricted in their choice of mechanism: they only need to guard against the buyer’s deviations towards any *single* other allocation. This allows the seller to utilize mechanisms where the buyer can only purchase a single mechanism entry. These *buy-one* mechanisms can be heavily tailored to the buyer’s distribution, often offering comparable allocations for widely different prices. Consider the following example.

Example 1. *A buyer walks into a coffee shop. They are equally likely to have one of three valuations over a cup of coffee and a bagel: either the buyer has value \$2 for the cup of coffee and \$0 for the bagel, \$0 for the cup of coffee and \$4 for the bagel, or \$4 for the cup of coffee and \$6 for the bagel (and \$10 for the combination of a cup of coffee and a bagel). The optimal mechanism in this example is as follows: the seller will offer the cup of coffee at \$2, the bagel at \$4 and the combination of a cup of coffee and a bagel at \$8. In this example, the optimal mechanism is buy-one incentive-compatible. The buyer with non-zero valuations for both items (weakly) prefers buying the combination at \$8 to buying exactly one of the items separately. The mechanism, however, is not buy-many incentive-compatible: when the buyer has non-zero value for both items, they would prefer to visit the coffee shop twice and buy the items separately for a combined price of \$6. This achieves the same allocation at a cheaper price.*

The example above highlights two problems with the classical buy-one model. The first is that no high-valued customer would pay \$8 for the combination of coffee and a bagel. They would buy one item, queue in line again, and buy the other. This in turn creates the second problem: the revenue of the optimal buy-one mechanism overshoots the real-world revenue the seller would experience. The buy-one mechanism would net the seller an expected revenue of $\$(2+4+8)/3 = \$4\frac{2}{3}$. In reality, because no buyer would pay \$8 for the combination of items, the seller would experience expected revenue $\$(2 + 4 + 6)/3 = \4 .

While this example is simple, the pathological constructions of [19, 26] do significantly wilder things such as offering similar randomized allocations for astronomically different prices. For instance, a buy-one mechanism may offer randomized allocation $\vec{q}$ for price p , and randomized allocation $\vec{q} + \vec{\varepsilon}$ for price $4p$. Just like in Example 1, a buyer would prefer to buy the cheaper option two, three or even four times rather than the carefully tailored, more expensive option. However, if they are forced to buy exactly one option, these pathological constructions ensure high-valued buyers will marginally prefer buying one copy of the expensive item to buying one copy of the cheaper one. Thus part of the reason why positive results in multi-item auctions (especially for correlated items) are rare is because the “optimal” buy-one mechanism is unrealistic and not implementable in a world where buyers may, reasonably, interact with the seller multiple times. It is this lack of consideration on the seller’s choice of mechanism that allows for “infinite” revenue auctions.

¹By “simple” mechanisms, we mean mechanisms of size polynomial in the number of items.

One natural way to overcome this problem is to allow the buyer to purchase multiple menu entries. *Buy-many* mechanisms, introduced more than ten years ago in [3, 4], are mechanisms where the buyer may purchase *any* multi-set of menu entries, including sets of unbounded size. This significantly restricts the seller’s choice of mechanism: buy-many mechanisms are always buy-one incentive compatible but the converse is not true. A simple way to see this is that the prices in (deterministic) buy-many mechanisms are always sub-additive, meaning that for any two disjoint sets of items S, T , $p(S) + p(T) \geq p(S \cup T)$. Buy-one mechanisms, like that of Example 1, need not satisfy this property, making them less appealing for real-world applications.

The work of [3, 4] already exhibits how buy-many mechanisms overcome the revenue gap problem: they showed that a popular benchmark, known as item-pricing, could recover an $O(\log n)$ factor of the revenue attained by the optimal buy-many mechanism for the case of a single, unit-demand buyer. This was later extended to arbitrary valuations by [12], while preserving the approximation factor. Key to these results is that by sufficiently restricting the seller’s choice of mechanisms, the optimal revenue drops from *unbounded* in the buy-one case to *finite* in the buy-many case, allowing for simple mechanisms like item-pricing to approximate the optimal buy-many revenue.

One question left unaddressed by these works is *how much* we need to restrict the seller’s choice of mechanisms so that the optimal revenue is finite. While in Example 1 it was reasonable to assume the buyer would purchase a single item, re-queue and purchase the other item, it would not be reasonable to assume the buyer would be willing to re-queue *any* number of times. At some point, the buyer will get tired. This means the buyer’s threat of interacting with the mechanism any number of times is not a credible threat to the seller. Alternatively, we can think of buy-many mechanisms as giving *too* much power to the buyer just like buy-one mechanisms give *too* much power to the seller.

In order to answer the question outlined we need a more fine-grained family of mechanisms that smoothly interpolates between buy-one and buy-many mechanisms. For this purpose we introduce buy- k mechanisms, a parametric family of mechanisms where the buyer is allowed to purchase any multi-set of at most k menu entries non-adaptively.² We say a mechanism is buy- k incentive-compatible if the buyer always prefers to buy a single menu entry rather than any multi-set of up to k menu entries. Let $\mathcal{B}_k(\mathcal{D})$ be the set of buy- k incentive-compatible mechanisms for a distribution $\mathcal{D}$ over n items, and let $\text{Buy}^k \text{Rev}(\mathcal{D}) = \max_{\mathcal{M} \in \mathcal{B}_k(\mathcal{D})} \text{Rev}(\mathcal{D}, \mathcal{M})$ be the optimal revenue attainable by a buy- k incentive-compatible mechanism. A simple observation, stated below and whose proof we defer to later in the paper, shows that as k increases, the revenue of the seller weakly decreases.

Observation 1.

$$\text{Buy}^1 \text{Rev}(\mathcal{D}) \geq \text{Buy}^2 \text{Rev}(\mathcal{D}) \geq \dots \geq \text{BuyManyRev}(\mathcal{D}) \geq \text{BRev}(\mathcal{D}).$$

We first show that k , the number of times the buyer might interact with the mechanism, can play a role in the seller’s optimal revenue. We prove there exists a support-size 3, correlated distribution $\mathcal{D}$ over two items for which $\text{Buy}^1 \text{Rev}(\mathcal{D}) > \text{Buy}^2 \text{Rev}(\mathcal{D}) > \text{Buy}^3 \text{Rev}(\mathcal{D}) > \text{Buy}^4 \text{Rev}(\mathcal{D})$. This example proves a strict separation between the class of buy-one and buy-2 mechanisms, and by Observation 1, between the class of buy-2 mechanisms and buy-many mechanisms. This reinforces the idea that buy- k mechanisms interpolate between buy-one and buy-many mechanisms, and thus merit a study of their own.

²In other words, the buyer first chooses any multi-set of up to k menu options and only after they commit any randomized allocations are decided. Our results will hold even if the buyer is allowed to adaptively choose the menu entries. See Appendix A for a more detailed discussion.

Proposition 1. *There exists a distribution $\mathcal{D}$ over two items such that*

$$\text{Buy}^1\text{Rev}(\mathcal{D}) > \text{Buy}^2\text{Rev}(\mathcal{D}) > \text{Buy}^3\text{Rev}(\mathcal{D}) > \text{Buy}^4\text{Rev}(\mathcal{D}).$$

We conjecture in fact that for the distribution from Proposition 1, the seller’s revenue strictly decreases as k increases. In other words, we conjecture there exists a simple distribution that can witness separations between the classes of buy- k and buy- $(k + 1)$ mechanisms for all $k \geq 1$. We discuss this conjecture in Appendix C.

Conjecture 1. *There exists a distribution $\mathcal{D}$ over two items such that for all $k \geq 1$,*

$$\text{Buy}^k\text{Rev}(\mathcal{D}) > \text{Buy}^{k+1}\text{Rev}(\mathcal{D}).$$

After proving that the class of buy- k mechanisms is distinct from the previously studied classes of buy-one and buy-many mechanisms, the next natural question is to understand their approximation guarantees with respect to simple mechanisms. Our measure of simplicity for a mechanism $\mathcal{M}$ will be its *menu complexity* or the number of menu entries $|\mathcal{M}|$ the mechanism offers. Under this lens, broadly speaking, we think of “simple” mechanisms as those that have polynomial menu complexity and “complex” mechanisms as those that have super-polynomial menu complexity. For example, any mechanism which only offers the grand bundle of all items for a fixed price has menu complexity 1. This family of mechanisms is so important that the revenue of the optimal grand bundling mechanism, $\text{BRev}(\cdot)$ (henceforth bundling), is often a benchmark of interest.³ Therefore, our main question of interest is the following.

Question 1. *Given integers n, k , when does $f(n, k) \cdot \text{BRev}(\mathcal{D}) \geq \text{Buy}^k\text{Rev}(\mathcal{D})$ hold for all distributions $\mathcal{D}$ over n items, for some function $f(n, k)$?*

1.1 Our Contributions

Our main result shows that restricting the seller to the class of buy- n incentive-compatible mechanisms suffices to get around pathological constructions for two or more items (like e.g., [19, 26]). These works show that there are distributions over just two items for which no “simple” mechanism could approximate the revenue of the optimal buy-one mechanism, or in the language of Open Question 1, that no such function $f(n, 1)$ exists for $n \geq 2$. We show that when facing a single, additive buyer, the revenue from optimally pricing the bundle of items, $\text{BRev}(\mathcal{D})$, recovers a polynomial fraction of the optimal buy- n revenue.

Theorem 1. *For any distribution $\mathcal{D}$ over n items for a single, additive buyer, it holds that*

$$O(n^2) \cdot \text{BRev}(\mathcal{D}) \geq \text{Buy}^n\text{Rev}(\mathcal{D}).$$

The proof of 1 relies on the identification of a measure, $\text{MenuGap}^k(\cdot, \cdot)$, whose formal definition we defer to Section 2. This quantity is the generalization to buy- k mechanisms of $\text{MenuGap}(\cdot, \cdot)$ introduced by previous work for buy-one mechanisms (see [19, 26]). In those works, $\text{MenuGap}(\cdot, \cdot)$ was used to construct distributions whose optimal revenue was hard to approximate. In contrast, our work is the first to show that this framework can be used to prove approximation guarantees

³Bundling is arguably one of the simplest mechanisms.

instead. In fact, our techniques can be extended to also show similar results for the case of arbitrary monotone valuation functions,⁴ (albeit with a significant loss in the approximation factor).

Theorem 2. *For any distribution $\mathcal{D}$ over n items, for a single buyer with a monotone valuation, it holds that*

$$O(n^2 \cdot 2^n) \cdot \text{BRev}(\mathcal{D}) \geq \text{Buy}^n \text{Rev}(\mathcal{D}).$$

We consider Theorem 2 as a result validating the robustness of the framework we introduce as a proof technique for approximation algorithms for revenue maximization. The experienced reader will recall that, historically, approximation algorithms for different valuations classes used tools specific to the valuations themselves (e.g., [7, 8] for unit demand, [28, 17, 2] for additive valuations and so on). It wasn't until the work of [5] that a unifying framework was developed to reprove (or even improve) such results. Thus, we interpret Theorem 2 additionally as proof that the framework we develop is robust enough to handle general valuation classes and are optimistic that results similar to Theorem 1 can be proved via our framework for other valuation classes.

The first piece of the proofs for Theorems 1, 2 is identical. We show that there exists an appropriate choice of inputs (X, Q) such that $\text{MenuGap}^k(X, Q)$ upper bounds the ratio between the optimal buy- k revenue and the revenue achieved by bundling, up to some $O(k)$ factor. We again defer a technical definition of $\text{MenuGap}^k(\cdot, \cdot)$ until later. For the purposes of this exposition, it suffices to say that given two sequences of vectors (X, Q) , $\text{MenuGap}^k(X, Q)$ captures some geometric property of the input pairs of sequences. Thus, more precisely, in the first step of the proof we show that for any distribution $\mathcal{D}$ and any buy- k incentive-compatible mechanism $\mathcal{M}$ for $\mathcal{D}$, there exists a cleverly chosen set of valuations X in the support of the distribution $\mathcal{D}$ together with their corresponding allocations Q under $\mathcal{M}$ whose “geometric property” witnesses an upper bound to the revenue revenue that $\mathcal{M}$ achieves on $\mathcal{D}$, up to a factor of $O(k)$. In particular, this is also true about the revenue-optimal buy- k incentive-compatible mechanism.

The second step of the proof upper bounds $\text{MenuGap}^n(X, Q)$ itself by n for the case of an additive buyer (resp. by $n \cdot 2^n$ for the case of a monotone buyer) for *all* input pairs (X, Q) . It is worth noting that for the case of an additive buyer, this second step is tight. This implies that our approach of bounding revenue gaps via $\text{MenuGap}^k(X, Q)$ cannot give a sublinear approximation. However, this is not a fault of our techniques. In Appendix B we provide a simple proof that there exist distributions $\mathcal{D}$ for which $\text{BRev}(\mathcal{D}) \leq \text{Buy}^k \text{Rev}(\mathcal{D}) / O(n)$ for any k . In other words, $\text{BRev}(\mathcal{D})$ can not give a sublinear (in the number of items n) approximation to $\text{Buy}^k \text{Rev}(\mathcal{D})$, for any k .

There are two subtle implications of these results. The first is that for all n -dimensional distributions $\mathcal{D}$, $\text{Buy}^n \text{Rev}(\mathcal{D})$ is *finite* whenever $\text{BRev}(\mathcal{D})$ is finite. This stands in contrast to the buy-one case where even for just $n = 2$ items, there exist $\mathcal{D}$ such that $\text{Buy}^1 \text{Rev}(\mathcal{D}) > \infty$ but $\text{BRev}(\mathcal{D}) = O(1)$. The second is that since $\text{Buy}^k \text{Rev}(\mathcal{D}) \geq \text{Buy}^{k'} \text{Rev}(\mathcal{D})$ whenever $k < k'$ (due to Observation 1), then Theorems 1, 2 in fact also answer Question 1 for the case when $n \leq k$.

The next goal is to answer Question 1 for the case when $1 < k < n$. We make progress by proving a strong lower bound for the case $k \leq n^{1/2-\varepsilon}$. We show that there exist distributions for which there is an exponential revenue gap when $k \leq n^{1/2-\varepsilon}$. This is captured by Theorem 3.

Theorem 3. *If $k \leq n^{1/2-\varepsilon}$ for some $\varepsilon > 0$, then there exists an additive valuation function $\mathcal{D}$ over n items such that for a single buyer*

⁴A valuation x function is monotone if $x(S) \geq x(T)$ whenever $T \subseteq S$. In other words, getting more items never decreases the buyer's value.

$$\frac{\text{Buy}^k \text{Rev}(\mathcal{D})}{\text{BRev}(\mathcal{D})} \geq \frac{\exp(\Omega(n^\epsilon))}{2n^2}.$$

The experienced reader will observe that Theorem 3 says something strong about the revenue guarantees that simple mechanisms can obtain. Due to a Corollary from [19] which says that bundling always recovers a $1/|\mathcal{M}|$ fraction of the revenue of any mechanism $\mathcal{M}$, the revenue of any mechanism $\mathcal{M}$ of size $\text{POLY}(n)$ cannot exceed $\text{POLY}(n) \cdot \text{BRev}(\mathcal{D})$. Thus, Theorem 3 implies that for the instance that witnesses its proof, no mechanism of size polynomial in the number of items can obtain a sub-exponential approximation. The proof of Theorem 3 will, unsurprisingly, borrow ideas from [19, 4]. Interestingly the buy- k mechanism used in the lower bound instance will be *deterministic*, in part because the construction of the instance itself makes use of discrete combinatorial objects known as cover-free sets. This implies that the lower bounds hold even for the optimal deterministic buy- k mechanism, a much weaker benchmark for comparison. To put this in context, [18] showed that there exist distributions $\mathcal{D}$ for which $\text{BRev}(\mathcal{D}) \leq \frac{2^n}{n} \text{DetRev}(\mathcal{D})$, where $\text{DetRev}(\mathcal{D})$ is the revenue of the optimal buy-one deterministic mechanism. Thus, Theorem 3 can be seen as an extension of this result, proving that even weakening the seller’s benchmark to the optimal buy- k deterministic mechanism (for $k \leq n^{1/2-\epsilon}$) does not significantly improve the worst-case approximation with respect to bundling.

While our model and results are written for the case of a non-adaptive buyer, a simple argument will allow us to translate both our upper bounds (Theorems 1, 2) and our lower bounds (Theorem 3) to the case of adaptive buyers. We defer this discussion to Appendix A.

Our Techniques. The main conceptual contribution of this work is the introduction of a new class of mechanisms, buy- k mechanisms, which interpolate between buy-one and buy-many mechanisms. The main technical contribution of our work is a novel framework for proving approximation results for multi-item mechanism design under arbitrary distributions. We generalize measures meant for the buy-one setting from [19, 26] to the buy- k setting. Similar to [26], we prove that this measure upper bounds the revenue gap between the revenue-optimal mechanism (in some class of mechanisms) and bundling. Unlike [26], we are able to show a *finite* upper bound for this measure in the case of buy- n mechanisms, yielding a finite approximation result.

1.2 Related Work

Buy-many mechanisms have been proposed more than ten years ago, as early as [3, 4]. Results from a recent line of work [12, 14, 13, 11] make the case to further the study of buy-many mechanisms. For instance, [14] show that buy-many mechanisms satisfy some form of revenue monotonicity, an intuitive property that does not hold in the case of buy-one mechanisms [20, 25]. In addition, as mentioned earlier in the introduction, [12] show that item-pricing recovers a $O(\log n)$ factor of the optimal buy-many revenue. Combining a Corollary from [19] with the main result of [12] shows that bundling recovers at least a $O(n \log n)$ fraction of the optimal buy-many revenue. Our results have a worse approximation factor because the benchmark is stronger (and this is proved formally in Observation 1). Thus our work deepens the study of buy-many mechanisms by introducing more fine-grained classes of mechanisms. We believe our results strengthen the case for the study of not only buy-many mechanisms, but also fine-grained buy-many mechanisms.

The work of [12] also gave strong lower bounds for the *description complexity*, a measure that lower bounds the menu complexity of a mechanism. In particular, they showed that no mechanism

with sub-exponential description complexity could get an $o(\log n)$ approximation to the optimal buy-many revenue, even for additive buyers. In follow up work, [14] extended the lower bound to the larger class of fractionally sub-additive (or XOS) valuations.

A prolific line of work assumes that the underlying distribution of values $\mathcal{D}$ is a product distribution. Under this assumption, it is known that mechanisms with low menu complexity can achieve constant-factor approximations to the optimal revenue for sub-additive valuations (see e.g., [23, 17, 28, 7, 9, 2, 6, 27, 8, 10, 1], among others), effectively circumventing the pathological constructions of [19]. Some recent results even show strong positive results for arbitrary approximation schemes. For instance, [1] show that for any product distribution $\mathcal{D}$, there exists a mechanism with finite menu complexity that recovers a $(1 - \varepsilon)$ approximation to the optimal revenue when selling to an additive buyer. More recently, [22] give a quasi-polynomial approximation scheme for revenue maximization for a single, unit-demand buyer interested in n independent items. Notwithstanding the significant contributions of these works, the question of revenue-maximization under arbitrary distributions remained unaddressed.

Finally, work of [25] provides yet another way to circumvent the pathological constructions of [19]. In their work, the authors borrow ideas from the celebrated smoothed-analysis framework and initiate the study of beyond worst-case revenue maximization. Their results show that, under some smoothing models, simple mechanisms can approximate optimal ones.

Organization. In Section 2, we present formal definitions for the objects of our interest as well as for the relevant benchmarks we use. Section 3 contains the proofs of our upper bounds in Theorems 1, 2. Section 4 contains the proof of our lower bound in Theorem 3. We conclude in Section 5 and outline questions for future work. Appendix A contains the discussion for the case of adaptive buyers. Appendix B proves that Theorem 1 can not be improved to provide a sub-linear approximation. Appendix C presents the proof of Proposition 1 and presents a simple candidate distribution to prove Conjecture 1.

2 Preliminaries and Notation

We consider the case of a single buyer interested in n items from a single revenue-maximizing seller. We assume the buyer is utility-maximizing and risk-neutral. The buyer draws their valuation x from a known, possibly correlated n -dimensional distribution $\mathcal{D}$, with support set $\mathcal{X} = \text{SUPP}(\mathcal{D})$ and probability density function $f_{\mathcal{D}}(\cdot)$ (and the index may be omitted when clear from context). For all our results, we will assume that the buyer’s valuation over the items is monotone, i.e., whenever $S \subseteq T$, $x(S) \leq x(T)$. For our main result, Theorem 1, we will additionally assume that the buyer is additive across the items, i.e., for any subset of items $S \subseteq [n]$, $x(S) = \sum_{i \in S} x_i$. Given a possibly randomized allocation of items $\vec{q} \in [0, 1]^n$, we use $x(\vec{q})$ to denote the buyer’s expected value for the realized set of items. For an additive buyer, $x(\vec{q}) = \sum_{i=1}^n x_i \cdot q_i$. For a monotone buyer, $x(\vec{q}) = \sum_{S \subseteq [n]} x(S) \cdot \Pr(\vec{q}, S)$, where $\Pr(\vec{q}, S) = \prod_{j \in S} \vec{q}_j \prod_{j \notin S} (1 - \vec{q}_j)$ is the probability that the buyer gets exactly the items in set S from randomized allocation $\vec{q}$. Let $\Lambda = \{\vec{q}_1, \vec{q}_2, \dots, \vec{q}_k, \dots\}$ be a multi-set of allocations (of possibly unbounded size). We denote by $\text{Lot}(\Lambda) \in [0, 1]^n$ (read “lottery”) the expected allocation that results from being allocated every $\vec{q}_i \in \Lambda$ independently and at once, i.e., for all $j \in [n]$, $\text{Lot}(\Lambda)_j = 1 - \prod_{i=1}^{\infty} (1 - q_{ij})$.⁵

⁵The careful reader might wonder what would happen if instead of buying all their menu options at once, the buyer was allowed to do so *adaptively* (as opposed to the model presented here which is non-adaptive). We discuss this in Appendix A, but the main takeaway is that our upper bounds also hold for adaptive buyers.

A mechanism $\mathcal{M} = (p, q)$ is defined by a pair of functions $p : \mathcal{X} \rightarrow \mathcal{R}_{\geq 0}$, $q : \mathcal{X} \rightarrow [0, 1]^n$ known as the pricing and allocation functions, respectively. For a fixed integer k we say that a mechanism $\mathcal{M}$ is *buy- k incentive-compatible* if for every valuation $\vec{x} \in \mathcal{X}$ it is in the buyer's best interest to purchase a single option from the mechanism rather than any combination of up to k menu options. In other words, there exists some $(p, \vec{q}) \in \mathcal{M}$ such that $x(\vec{q}(x)) - p(x) \geq x(\text{Lot}(\Lambda)) - \sum_{i=1}^k p(\vec{q}_i)$ for *any* possible multi-set of menu options Λ of size at most k . Thus, setting $k = 1$ recovers the standard notion of (buy-one) incentive-compatible mechanisms, and as $k \rightarrow \infty$ it recovers the existing definition of buy-many incentive-compatible mechanisms.

2.1 Benchmarks of Interest

We now formally define some of the benchmarks that will be used throughout this paper. For a given (arbitrarily correlated) distribution $\mathcal{D}$, let

- $\text{BRev}(\mathcal{D})$ be the revenue of the mechanism which sells the grand bundle for its optimal price. Namely, $\text{BRev}(\mathcal{D}) = \max_p p \cdot \Pr_{\vec{x} \sim \mathcal{D}}(\sum x_i \geq p)$,
- $\text{Rev}(\mathcal{D})$ be the revenue of the optimal buy-one incentive-compatible mechanism,
- $\text{Rev}(\mathcal{D}, \mathcal{M})$ be the revenue of mechanism $\mathcal{M}$ when the buyer is allowed to buy up to 1 menu entry from $\mathcal{M}$,
- $\text{Buy}^k \text{Rev}(\mathcal{D})$ be the revenue of the optimal buy- k incentive-compatible mechanism,
- $\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M})$ be the revenue of (not necessarily buy- k incentive-compatible) mechanism $\mathcal{M}$ when the buyer is allowed to buy up to k menu entries from $\mathcal{M}$,
- $\text{BuyManyRev}(\mathcal{D})$ be the revenue of the optimal buy-many incentive-compatible mechanism.

Proof of Claim:1. If a mechanism $\mathcal{M}$ is buy- k incentive-compatible for some k , it is also buy- k' incentive-compatible for all $k' \leq k$. This follows from the fact that the buyer can always buy the empty lottery $k - k'$ times and the mechanism must guard against such deviations. Thus, the best buy- k mechanism can perform no better than the best buy- $(k - 1)$ mechanism, proving the claim. The last inequality follows from the fact that bundling is a buy-many incentive compatible mechanism. $\square$

2.2 Menu Gaps: An Intermediary Measure

We now present the definition of $\text{gap}_i^k(X, Q)$ and $\text{MenuGap}^k(X, Q)$, the quantities that will serve as intermediaries in proving Theorems 1, 2.

Definition 1. Let $X = \{x_i\}_{i=1}^N$ be a sequence of monotone valuation functions and $Q = \{\vec{q}_i\}_{i=0}^N \in [0, 1]^n$ be a sequence of vectors with $\vec{q}_0 = \vec{0}^n$. Then

$$\text{gap}_i^k(X, Q) = \min_{j_1, j_2, \dots, j_k < i} x_i(\vec{q}_i) - x_i(\text{Lot}(\vec{q}_{j_1}, \vec{q}_{j_2}, \dots, \vec{q}_{j_k})). \quad (1)$$

Furthermore, define

$$\text{MenuGap}^k(X, Q) = \sum_{i=1}^N \text{gap}_i^k(X, Q) / x_i([n]), \quad (2)$$

where $x_i([n])$ is the valuation of a buyer of type x_i for the grand bundle of items.

In particular, if $X = \{x_i\}_{i=1}^N$ are additive valuations, then they can be represented as vectors $X = \{\vec{x}_i\}_{i=1}^N \in \mathbb{R}_{\geq 0}^n$. One can check that the following definition is a special case of Definition 1.

Definition 2. Let $X = \{\vec{x}_i\}_{i=1}^N \in \mathbb{R}_{\geq 0}^n, Q = \{\vec{q}_i\}_{i=0}^N \in [0, 1]^n$ be sequences of vectors with $\vec{q}_0 = \vec{0}^n$. Then

$$\text{gap}_i^k(X, Q) = \min_{j_1, j_2, \dots, j_k < i} \vec{x}_i \cdot (\vec{q}_i - \vec{\text{Lot}}(\vec{q}_{j_1}, \vec{q}_{j_2}, \dots, \vec{q}_{j_k})). \quad (3)$$

Furthermore, define

$$\text{MenuGap}^k(X, Q) = \sum_{i=1}^N \text{gap}_i^k(X, Q) / \|\vec{x}_i\|_1, \quad (4)$$

These measures are generalizations of similar notions introduced in [19] and further developed by [26]. For the case where $k = 1$, we exactly recover these earlier definitions. In Definition 1, it is useful to think of the first sequence of vectors X as possible valuation vectors and the sequence of vectors Q as possible allocation vectors of a mechanism, with the built-in option of not participating. Thus, one way to interpret Equation 1 is to think of $p_i = \text{gap}_i^1(X, Q)$ as the largest price a seller can post on menu entry $(p_i, \vec{q}_i)$ so that a buyer with valuation x_i will prefer that single menu entry to any subset of at most k “previous” options for *free*.

2.3 Some Useful Properties

We prove a simple, useful property of the $\vec{\text{Lot}}(\Lambda)$ function. Namely, that if $\Lambda = \{\vec{q}_1, \vec{q}_2, \dots, \vec{q}_k\}$, then $\vec{\text{Lot}}(\Lambda)$ dominates the vector consisting of coordinate-wise max entries of the vectors in Λ .

Claim 4. If $\Lambda = \{\vec{q}_1, \vec{q}_2, \dots, \vec{q}_k\}, \vec{q}_i \in [0, 1]^n$ for all i , then $\vec{\text{Lot}}(\Lambda)_j \geq \max_{i \in [k]} \{q_{ij}\}$ for all $j \in [n]$.

Proof. Fix $j \in [n]$. Assume wlog $q_{1j} = \max_{i \in [k]} q_{ij}$. It is easy to see that $(1 - q_{1j}) \cdot (1 - \prod_{i=2}^k (1 - q_{ij})) \geq 0$ since each term on the left is non-negative. Expanding it we get that $1 - q_{1j} - \prod_{i=2}^k (1 - q_{ij}) \geq 0$. Rewriting gives $\vec{\text{Lot}}(\Lambda)_j \geq q_{1j}$. $\square$

We also prove a simple, useful property of the $\text{MenuGap}^k(X, Q)$ function. Namely, that it is without loss of generality to remove points whose contributions to the sum are non-positive.

Claim 5. Let X, Q be sequences as defined in Definition 1, and $X' \subseteq X, Q' \subseteq Q$ be the sub-sequences that result from removing any pair of points $(\vec{x}_i, \vec{q}_i)$ whose $\text{gap}_i^k(X, Q) \leq 0$. Then

$$\text{MenuGap}^k(X, Q) \leq \text{MenuGap}^k(X', Q').$$

Proof. Consider the earliest integer i such that $\text{gap}_i^k(X, Q) \leq 0$. Since it is non-positive, removing $(x_i, \vec{q}_i)$ from (X, Q) will weakly increase the sum of the gaps up to i . Moreover, if $\vec{q}_i$ was helping set the gap for some later $(x_j, \vec{q}_j)$, then $\text{gap}_j^k(X', Q') \geq \text{gap}_j^k(X, Q)$ since by removing $\vec{q}_i$ we are reducing the number of earlier points to compare to. Therefore, removing any point with negative gap can only weakly increase the menu gap of the resulting subsequence. $\square$

3 Bundling Approximates the Optimal Buy-N Revenue

As highlighted in the introduction of the paper, the first part of the proofs of Theorems 1, 2 will be via the surrogate quantity, $\text{MenuGap}^k(X, Q)$. We will first show that for any distribution over n items $\mathcal{D}$, there exists two sequences of points (X, Q) such that $\text{MenuGap}^k(X, Q)$ upper bounds the ratio between the revenue-optimal buy- k mechanism for $\mathcal{D}$ and the revenue from bundling, up to a $O(k)$ factor:

Lemma 6. *For any distribution of monotone valuations $\mathcal{D}$ over n items and any buy- k incentive compatible mechanism $\mathcal{M}$ for $\mathcal{D}$, there exists a sequence of valuations $X = \{x_i\}_{i=1} \subseteq \mathcal{X}$, and a sequence of allocations $Q = \{\vec{q}_i\}_{i=0} \subseteq \mathcal{M}$ (starting with $\vec{q}_0 = (0, \dots, 0)$) such that*

$$\text{MenuGap}^k(X, Q) \geq \frac{\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M})}{9k \cdot \text{BRev}(\mathcal{D})}.$$

Next we will show that this quantity itself is upper bounded for *all* pairs of sequences (X, Q) . In particular, when the buyer has additive valuations, we show that the quantity is upper bounded by n . The proof of Theorem 1 then follows directly.

Lemma 7. *For all sequences X, Q as defined in Definition 2, coming from an additive valuation function, it holds that $\text{MenuGap}^n(X, Q) \leq n$.*

Theorem 1. *For any distribution $\mathcal{D}$ over n items for a single, additive buyer, it holds that*

$$O(n^2) \cdot \text{BRev}(\mathcal{D}) \geq \text{Buy}^n \text{Rev}(\mathcal{D}).$$

Proof of Theorem 1. Follows directly from Lemma 6 by choosing $\mathcal{M}$ to be the revenue-optimal buy- k incentive-compatible mechanism and Lemma 7 (setting $k = n$). $\square$

The proof of Theorem 2 is similar, but only a weaker version of Lemma 7 can be proved:

Theorem 2. *For any distribution $\mathcal{D}$ over n items, for a single buyer with a monotone valuation, it holds that*

$$O(n^2 \cdot 2^n) \cdot \text{BRev}(\mathcal{D}) \geq \text{Buy}^n \text{Rev}(\mathcal{D}).$$

Lemma 8. *For all sequences X, Q as defined in Definition 1, coming from a monotone valuation function, it holds that $\text{MenuGap}^n(X, Q) \leq n \cdot 2^n$.*

Proof of Theorem 2. Follows directly from Lemma 6 by choosing $\mathcal{M}$ to be the revenue-optimal buy- k incentive-compatible mechanism and Lemma 8 (setting $k = n$). $\square$

Note that in Lemma 6, the number of times the buyer can interact with the mechanism, k , may be different than the number of items n for sale. However, we are only able to prove Lemma 8 for the case where $k \geq n$. We hope future work can address the question of what happens when $k < n$.

3.1 Menu Gap Approximately Upper Bounds Revenue Gap: Proof of Lemma 6

The proof of Lemma 6 is split into two parts. In the first part, we will any buy- k incentive-compatible mechanism for $\mathcal{D}$ and massage it down to a sub-menu of interest whose revenue remains close to the original one. The sub-menu itself may not be buy- k incentive-compatible. However, the key to approximately preserving the revenue will be in just removing entries from the original mechanism and not modifying existing ones. In the second part, we will show how to use an appropriate sub-menu in order to construct the desired sequence of points. The proof of Lemma 6 is inspired on a similar construction of [26].

3.1.1 Finding a Sub-menu of Interest

Let $\mathcal{M} = \{(p_i, \vec{q}_i)\}_{i=1}$ be a buy- k incentive compatible mechanism for distribution $\mathcal{D}$, where $(p_i, \vec{q}_i)$ denotes the price and expected allocation of the i -th option of the menu.

Claim 9. *Let $\mathcal{M}$ be a buy- k incentive-compatible mechanism for $\mathcal{D}$, $\mathcal{M}_c \subseteq \mathcal{M}$ be the sub-menu of $\mathcal{M}$ that only offers options of price at least c . Then $\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}_c) \geq \text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}) - c$.*

Proof. If a buyer with valuation x chose a menu entry $(p, \vec{q})$ from the original menu $\mathcal{M}$ with $p \geq c$, they will purchase the same menu entry in $\mathcal{M}_c$ since $(p, \vec{q})$ was utility-maximizing and no new menu entries were introduced. If $p < c$, it is possible that the buyer would purchase some other option $(p', \vec{q}')$ (or combination of options). Regardless, the loss in revenue from that buyer is bounded by $cf(x)$. Let S_c be the set of valuation vectors that preferred a menu entry in $\mathcal{M}$ priced at $p < c$. Then the total loss in revenue is at most $c \sum_{x \in S_c} f(x) \leq c$. $\square$

Claim 10. *Let $\mathcal{M}$ be a mechanism, and let $\mathcal{M}_1, \mathcal{M}_2 \subseteq \mathcal{M}$ be sub-menus of $\mathcal{M}$ defined as follows:*

- $\mathcal{M}_1$ has all options whose price $p_i \in \cup_{i=0}^{\infty} [c \cdot (2k)^{2i}, c \cdot (2k)^{2i+1})$.
- $\mathcal{M}_2$ has all options whose price $p_i \in \cup_{i=0}^{\infty} [c \cdot (2k)^{2i+1}, c \cdot (2k)^{2i+2})$.

Then $\max_{i=1,2} \text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}_i) \geq \text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M})/2$.

Proof. Because $\mathcal{M}_1 \cup \mathcal{M}_2 = \mathcal{M}$, observe that

$$\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}) \leq \text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}_1) + \text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}_2).$$

This is because any buyer with valuation x who purchases an option from $\mathcal{M}_1$ when presented the menu $\mathcal{M}$ will buy the same option when only presented $\mathcal{M}_1$. By a simple averaging argument, the better of the two menus must get revenue at least half of the revenue of the original menu. $\square$

Lemma 11. *There exists a menu $\mathcal{M}$ such that*

- All prices are at least c .
- All prices belong to the set of intervals $\cup_{i=0}^{\infty} [c \cdot (2k)^{2i+a}, c \cdot (2k)^{2i+a+1})$ for an $a \in \{0, 1\}$.
- $\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}) \geq \frac{\text{Buy}^k \text{Rev}(\mathcal{D}) - c}{2}$.

Proof. Take the initial buy- k incentive-compatible menu $\mathcal{M}$, apply Claim 9 to obtain a menu $\mathcal{M}'$ that satisfies the first bullet point. Take the menu $\mathcal{M}'$ and apply Claim 10 to obtain a menu $\mathcal{M}''$ that immediately satisfies the first and second bullet points. Finally, due to the revenue guarantees of Claims 9, 10, we have that $\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}'') \geq \frac{\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}')}{2} \geq \frac{\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}) - c}{2}$. $\square$

3.1.2 Construction of the Sequences X, Q

In this subsection we will show how to use the sub-menu found in the previous subsection to construct the sequences (X, Q) of interest who would witness

$$\text{MenuGap}^k(X, Q) \geq O(\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}) / \text{BRev}(\mathcal{D})).$$

Consider the menu $\mathcal{M}''$ from Lemma 11. Let $\mathcal{B}_j \subseteq \mathcal{M}''$ be the sub-menu that has all menu entries priced in $[c \cdot (2k)^{2j+a}, c \cdot (2k)^{2j+a+1})$ for the same $a \in \{0, 1\}$ from Lemma 11. We say a valuation $x \in \mathcal{B}_j$ if the menu option $(p(x), \vec{q}(x)) \in \mathcal{B}_j$. Let x_j be the valuation on $\mathcal{B}_j$ such that $x_j([n]) \leq (1 + \delta)x([n]) \forall x \in \mathcal{B}_j$. We call valuation x_j the *representative of bin $\mathcal{B}_j$* .

Claim 12. $\Pr(x \in \mathcal{B}_j) \leq \frac{\text{BRev}(\mathcal{D})(1+\delta)}{x_j([n])}$.

Proof. Consider the mechanism that sells the grand bundle at price $x_j([n])/(1 + \delta)$. Since any valuation on $\mathcal{B}_j$ has value at least that much for the grand bundle, the revenue of this menu is at least $\Pr(x \in \mathcal{B}_j)x_j([n])/(1 + \delta)$. But this is a grand bundling menu and is thus its revenue is at most $\text{BRev}(\mathcal{D})$. $\square$

Let (X, Q) be the sequence defined by the choice of x_j and their respective allocations in $\mathcal{M}$, $\vec{q}_j$.

Claim 13. $\frac{\text{gap}_j^k(X, Q)}{x_j([n])} \geq \frac{p_j}{2 \cdot x_j([n])} \geq \frac{\Pr(x \in \mathcal{B}_j)p_j}{2 \cdot \text{BRev}(\mathcal{D})(1+\delta)}$.

Proof. Because the initial mechanism $\mathcal{M}^*$ was buy- k incentive-compatible, we know that for any previous set of k options $\vec{q}_{j_1}, \dots, \vec{q}_{j_k}$,

$$x_j(\vec{q}_j) - p_j \geq x_j(\text{Lot}(\vec{q}_{j_1}, \dots, \vec{q}_{j_k})) - \sum_{i=1}^k p_{j_i}.$$

We can rewrite this as

$$\frac{\text{gap}_j^k(X, Q)}{x_j([n])} \geq \frac{p_j - \sum_{i=1}^k p_{j_i}}{x_j([n])}.$$

Recall by our choice of points and the fact that $j_i < j$, $p_j \geq 2k \cdot p_{j_i}$. Therefore, the right hand is at least $\frac{p_j}{2x_j([n])}$. The second inequality comes from Claim 12. $\square$

We are now ready to prove Lemma 6.

Proof of Lemma 6. Let us first observe that

$$\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}'') = \sum_j \sum_{x \in \mathcal{B}_j} p(x)f(x) \leq \sum_j \Pr(x \in \mathcal{B}_j)2kp_j. \quad (5)$$

Recall that the price any valuations $x \in \mathcal{B}_j$, its price $p(x)$ is no greater than $2kp_j$. Therefore, the inequality follows. Moreover, from Claim 13 we get that

$$\text{MenuGap}^k(X, Q) = \sum_j \frac{\text{gap}_j^k(X, Q)}{x_j([n])} \geq \sum_j \frac{\Pr(x \in \mathcal{B}_j)p_j}{2 \cdot \text{BRev}(\mathcal{D})(1 + \delta)}. \quad (6)$$

Applying Eq. 5 together with Lemma 11 we get that

$$\sum_j \Pr(x \in \mathcal{B}_j) p_j \geq \frac{\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}'')}{2k} \geq \frac{(\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}) - c)}{4k}. \quad (7)$$

Putting Eqs. 6, 7 we get that

$$\text{MenuGap}^k(X, Q) \geq \frac{(\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M}) - c)}{8k \cdot \text{BRev}(\mathcal{D})(1 + \delta)}. \quad (8)$$

Let $c = \text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M})/100$, $\delta = 1/100$ in Equation 8. Therefore,

$$\frac{99 \cdot \text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M})}{101 \cdot 8k \cdot \text{BRev}(\mathcal{D})} \geq \frac{\text{Buy}^k \text{Rev}(\mathcal{D}, \mathcal{M})}{9 \cdot k \cdot \text{BRev}(\mathcal{D})}.$$

□

3.2 Menu Gap is Finite when $k \geq n$: Proof of Lemmas 7, 8

The proof of Lemmas 7, 8 will be similar. We introduce some common notation to both Lemmas before proving each of them individually. Given a sequence of points Q , let Q_i be the sequence truncated at the i -th point, that is to say $Q_i = \{\vec{q}_0, \vec{q}_1, \dots, \vec{q}_i\}$. Let $\vec{m}_i = (\max_{\vec{q}_i \in Q_i} \{\vec{q}_{i,1}\}, \dots, \max_{\vec{q}_i \in Q_i} \{\vec{q}_{i,n}\})$ be the n -dimensional vector whose entries are the largest coordinates among the points in Q_i .

3.2.1 Proof Lemma 7

In this subsection we abuse the fact that for additive buyers, we can think of their valuation function x as an n -dimensional vector $\vec{x} = (x_1, \dots, x_n)$ and $x(\vec{q}) = \vec{x} \cdot \vec{q}$.

Claim 14. For any i , $\text{gap}_i^n(X, Q)/\|\vec{x}_i\|_1 \leq \sum_{d=1}^n \max\{\vec{q}_{i,d} - \vec{m}_{i-1,d}, 0\}$.

Proof. Let $\vec{x}_i, \vec{q}_i$ be given. Since gap is defined to be the minimum over all pairs of previously placed points, we can just upper bound it by witnessing its value with earlier points. Let $i_1^*, i_2^*, \dots, i_n^*$ be the indices such that:

- $i_1^*, i_2^*, \dots, i_n^* < i$,
- $\vec{q}_{i_d^*, d} = \vec{m}_{i-1, d} \forall d \in [n]$.

That is to say, $\{i_d^*\}_{d=1}^n$ are the indices of the points that witness that $\vec{m}_{i-1}$ is indeed the coordinate-wise max of all points in Q_{i-1} . Then

$$\text{gap}_i^n(X, Q)/\|\vec{x}_i\|_1 \leq \frac{\vec{x}_i}{\|\vec{x}_i\|_1} \cdot (\vec{q}_i - \vec{\text{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*}, \dots, \vec{q}_{i_n^*})).$$

Recall that

$$\vec{\text{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*}, \dots, \vec{q}_{i_n^*})_d \geq \max\{\vec{q}_{i_1^*}, \vec{q}_{i_2^*}, \dots, \vec{q}_{i_n^*}\}_d \geq \vec{m}_{i-1, d}.$$

Therefore,

$$\vec{q}_{i, d} - \vec{\text{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*}, \dots, \vec{q}_{i_n^*})_d \leq \vec{q}_{i, d} - \vec{m}_{i-1, d},$$

for all $d \in [n]$. Therefore, for any choice of $\vec{x}_i$, it will be true that

$$\begin{aligned}
\text{gap}_i^n(X, Q) / \|\vec{x}_i\|_1 &\leq \frac{\vec{x}_i}{\|\vec{x}_i\|_1} \cdot (\vec{q}_i - \vec{\text{Lot}}(\vec{q}_{i_1}^*, \vec{q}_{i_2}^*, \dots, \vec{q}_{i_n}^*)) \\
&= \sum_{d=1}^n \frac{\vec{x}_{i,d}}{\|\vec{x}_i\|_1} (\vec{q}_{i,d} - \vec{\text{Lot}}(\vec{q}_{i_1}^*, \vec{q}_{i_2}^*, \dots, \vec{q}_{i_n}^*)_d) \\
&\leq \sum_{d=1}^n \frac{\vec{x}_{i,d}}{\|\vec{x}_i\|_1} (\vec{q}_{i,d} - \vec{m}_{i-1,d}) \\
&\leq \sum_{d=1}^n \max\{0, \vec{q}_{i,d} - \vec{m}_{i-1,d}\}
\end{aligned}$$

This proves the claim (Naturally, $\vec{x}_{i,d} \leq \|\vec{x}_i\|_1$ for all d). $\square$

Proof of Lemma 7. For any pair of sequences (X, Q) coming from an additive valuation function,

$$\begin{aligned}
\text{MenuGap}^n(X, Q) &= \sum_{i=1}^N \text{gap}_i^n(X, Q) / \|\vec{x}_i\|_1 \\
&\leq \sum_{d=1}^n \sum_{i=1}^N (\max\{\vec{q}_{i,d} - \vec{m}_{i-1,d}, 0\}) \\
&\leq \sum_{d=1}^n \sum_{i=1}^N (\max\{\vec{m}_{i,d} - \vec{m}_{i-1,d}, 0\}) \\
&\leq \sum_{d=1}^n \sum_{i=1}^N (\vec{m}_{i,d} - \vec{m}_{i-1,d}) \\
&\leq n.
\end{aligned}$$

The first inequality follows from Claim 14. For the second inequality, first note that for any fixed d , by definition $\vec{q}_{i,d} \leq \vec{m}_{i,d}$, with equality only if $\vec{q}_{i,d} \geq \vec{q}_{i',d}$ for all $i' < i$. But note also that by definition $\vec{m}_{i,d} - \vec{m}_{i-1,d} \geq 0$. Therefore $\max\{\vec{m}_{i,d} - \vec{m}_{i-1,d}, 0\} \leq \vec{m}_{i,d} - \vec{m}_{i-1,d}$. The last inequality follows from observing that the final sum across each coordinate telescopes. Since $\vec{q}_{i,d} \leq 1$, the sum is at most 1 per coordinate. $\square$

Observation 15. *Setting both sequences (X, Q) equal to the standard basis of $\mathbb{R}^n$ shows that Lemma 7 is tight.*

3.2.2 Proof of Lemma 8

In this subsection we assume the valuation function is monotone, i.e. $v(S) \geq v(T)$ whenever $T \subseteq S$. We will show that for such valuation classes, $\text{MenuGap}^n(X, Q) \leq 2^n \cdot n$ for all sequences (X, Q) . To prove this, we'll make use of the following lemma.

Claim 16. *For any monotone valuation v and two vectors $p, q \in [0, 1]^n$ with $p \leq q$ coordinate-wise (i.e., $p_i \leq q_i$ for all $i \in [n]$), we have $v(p) \leq v(q)$.*

Proof. To simplify the proof, observe that we only need to consider the case where p and q differ in exactly one coordinate (i.e., $p_i = q_i$ for $i \in [n-1]$ and $p_n < q_n$). We can then apply this result (up to) n times in order to conclude our desired result.

Since we can partition the powerset of $[n]$ into sets which contain n and sets which don't contain n , we can rewrite $v(\vec{p})$ as:

$$v(\vec{p}) = \sum_{A \subseteq [n]} \Pr(\vec{p}, A) \cdot v(A) = \sum_{B \subseteq [n-1]} \Pr(\vec{p}, B \cup \{n\}) \cdot v(B \cup \{n\}) + \sum_{B \subseteq [n-1]} \Pr(\vec{p}, B) \cdot v(B) \quad (9)$$

This implies that $v(\vec{q}) - v(\vec{p}) = A + B$, where A, B are defined as:

$$A = \sum_{B \subseteq [n-1]} [\Pr(\vec{q}, B \cup \{n\}) - \Pr(\vec{p}, B \cup \{n\})] \cdot v(B \cup \{n\}) \quad (10)$$

$$B = \sum_{B \subseteq [n-1]} [\Pr(\vec{q}, B) - \Pr(\vec{p}, B)] \cdot v(B) \quad (11)$$

Our goal is to show $A + B \geq 0$. To do this, let's first find explicit expressions of $[\Pr(\vec{q}, B \cup \{n\}) - \Pr(\vec{p}, B \cup \{n\})]$ and $[\Pr(\vec{q}, B) - \Pr(\vec{p}, B)]$. We have:

$$[\Pr(\vec{q}, B \cup \{n\}) - \Pr(\vec{p}, B \cup \{n\})] = (q_n - p_n) \prod_{i \in B} q_i \cdot \prod_{i \notin B \cup \{n\}} (1 - q_i) \quad (12)$$

and

$$[\Pr(\vec{q}, B) - \Pr(\vec{p}, B)] = [(1 - q_n) - (1 - p_n)] \prod_{i \in B} q_i \cdot \prod_{i \notin B \cup \{n\}} (1 - q_i) \quad (13)$$

by direct calculation (and noting that $p_i = q_i$ for all $i \in [n-1]$). Thus, we observe that the two probabilities are additive complements of each other. Hence, if we denote $p_B = [\Pr(\vec{q}, B \cup \{n\}) - \Pr(\vec{p}, B \cup \{n\})]$, we can write

$$v(\vec{q}) - v(\vec{p}) = \sum_{B \subseteq [n-1]} p_B \cdot [v(B \cup \{n\}) - v(B)]. \quad (14)$$

But we know the right hand side is at least 0 since $v(B \cup \{n\}) \geq v(B)$ by monotonicity. Hence, we can conclude $v(\vec{q}) - v(\vec{p}) \geq 0 \implies v(\vec{p}) \leq v(\vec{q})$. $\square$

We are now ready to prove Lemma 8.

Proof of Lemma 8. Recall that, for all i , we can choose $i_1^*, \dots, i_n^* < i$ such that $\vec{q}_{i_d^*, d} = \vec{m}_{i-1, d}$ for all $d \in [n]$. Thus, by monotonicity and Claim 16, we have $\max_{j_1, \dots, j_n < i} v(\text{Lot}(\vec{q}_{j_1}, \dots, \vec{q}_{j_n})) \geq v(\vec{m}_{i-1})$. Additionally, by definition of $\vec{m}_i$ and monotonicity (combined with Claim 16), we have $v(\vec{q}_i) \leq v(\vec{m}_i)$. Consequently, we can write

$$\frac{\text{gap}_i^k(X, Q)}{v_i([n])} = \min_{j_1, \dots, j_k < i} \frac{v_i(\vec{q}_i) - v_i(\text{Lot}(\vec{q}_{j_1}, \dots, \vec{q}_{j_k}))}{v_i([n])} \leq \frac{v_i(\vec{m}_i) - v_i(\vec{m}_{i-1})}{v_i([n])}$$

since v_i are assumed to be monotone. Writing out the definition of v_i , we know

$$\begin{aligned} \frac{v_i(\vec{m}_i) - v_i(\vec{m}_{i-1})}{v_i([n])} &= \sum_{A \in 2^{[n]}} \frac{v_i(A)}{v_i([n])} \cdot [\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)] \\ &\leq \sum_{A \in 2^{[n]}} |\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)| \end{aligned}$$

since $v_i(A) \leq v_i([n])$ by monotonicity. Combining this with above yields

$$\begin{aligned} \text{MenuGap}^n(X, Q) &= \sum_{i=1}^N \frac{\text{gap}_i^n(X, Q)}{v_i([n])} \\ &\leq \sum_{A \in 2^{[n]}} \sum_{i=1}^N |\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)| \end{aligned}$$

It remains to show $\sum_{i=1}^N |\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)| \leq n$ for all $A \in 2^{[n]}$. It would then follow directly that $\text{MenuGap}^n(X, Q) \leq 2^n \cdot n$. Fix a set A and consider each term $|\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)|$ individually.

Let $\vec{m}_i = \vec{m}^{(0)}, \vec{m}^{(1)}, \dots, \vec{m}^{(n)} = \vec{m}_{i-1}$ be a sequence of vectors where $\vec{m}^{(j)}$ matches $\vec{m}_i$ on the first j coordinates and matches $\vec{m}_{i-1}$ on the last $n - j$ coordinates. By the triangle inequality,

$$|\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)| \leq \sum_{j=1}^n |\Pr(\vec{m}^{(j)}, A) - \Pr(\vec{m}^{(j-1)}, A)|.$$

But by definition of $\Pr(\vec{q}, S)$, we know

$$|\Pr(\vec{m}^{(j)}, A) - \Pr(\vec{m}^{(j-1)}, A)| = \left| \prod_{\ell \in A} \vec{m}_\ell^{(j)} \prod_{\ell \notin A} (1 - \vec{m}_\ell^{(j)}) - \prod_{\ell \in A} \vec{m}_\ell^{(j-1)} \prod_{\ell \notin A} (1 - \vec{m}_\ell^{(j-1)}) \right| \leq \vec{m}_{i,j} - \vec{m}_{i-1,j}$$

since $\vec{m}^{(j)}$ and $\vec{m}^{(j-1)}$ differ only in the j^{th} coordinate and the remaining probabilities are at most one. Hence, combining with the above, we have

$$|\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)| \leq \sum_{j=1}^n \vec{m}_{i,j} - \vec{m}_{i-1,j}.$$

Summing over $i \in [N]$, we see that the sum telescopes:

$$\sum_{i=1}^N |\Pr(\vec{m}_i, A) - \Pr(\vec{m}_{i-1}, A)| \leq \sum_{j=1}^n \sum_{i=1}^N \vec{m}_{i,j} - \vec{m}_{i-1,j} \leq \sum_{j=1}^n \vec{m}_{N,j} \leq n,$$

since $\vec{m}_{N,j} \leq 1$ for all $j \in [n]$. □

4 Few Tickets Do Not Suffice: Proof of Theorem 3

In this section we show that if $k \leq n^{1/2-\varepsilon}$, then there exists a distribution $\mathcal{D}$ over n items for which there is an exponential gap in n (up to $\text{POLY}(n)$) between $\text{BRev}(\mathcal{D})$ and $\text{Buy}^k\text{Rev}(\mathcal{D})$, the revenue attained by the optimal buy- k mechanism.

Theorem 3. *If $k \leq n^{1/2-\varepsilon}$ for some $\varepsilon > 0$, then there exists an additive valuation function $\mathcal{D}$ over n items such that for a single buyer*

$$\frac{\text{Buy}^k\text{Rev}(\mathcal{D})}{\text{BRev}(\mathcal{D})} \geq \frac{\exp(\Omega(n^\varepsilon))}{2n^2}.$$

The proof of Theorem 3 will be broken down in three steps. Firstly, we will describe the pair of sequences (X^L, Q^L) that we use (Subsection 4.1.1). The construction will make use of a combinatorial Lemma about cover-free sets from [21]. Next, we will show that for that instance, $\text{MenuGap}^k(X^L, Q^L) \geq \frac{\exp(\Omega(n^\varepsilon))}{2n^2}$ when $k \leq n^{1/2-\varepsilon}$ (Lemma 20). In the final step, we will show how to construct a distribution $\mathcal{D}$ such that $\text{Buy}^k\text{Rev}(\mathcal{D})/\text{BRev}(\mathcal{D}) \geq \text{MenuGap}^k(X^L, Q^L)$ (Lemma 21). The proof of Lemma 21 will use ideas from [19].

Before we delve into the proof of Theorem 3, we analyze its implications for mechanisms with polynomial menu size. We invoke the following Corollary from [19].

Corollary 17 (Restated from [19]). *Consider any mechanism $\mathcal{M}$ with menu size M , then for any distribution $\mathcal{D} \in \mathbb{R}^n$*

$$M \cdot \text{BRev}(\mathcal{D}) \geq \text{Rev}(\mathcal{D}, \mathcal{M}).$$

This Corollary, combined with Theorem 3 imply the following Corollary.

Corollary 18. *Let $\mathcal{M}$ be a buy- k mechanism with menu size $M = \text{POLY}(n)$, then there exists a distribution $\mathcal{D} \in \mathbb{R}^n$ such that for any single, additive buyer*

$$\frac{\text{Buy}^k\text{Rev}(\mathcal{D})}{\text{Rev}(\mathcal{D}, \mathcal{M})} \geq \frac{\exp(\Omega(n^\varepsilon))}{\text{POLY}(n)}.$$

In other words, Corollary 18 rules out all polynomial-sized mechanisms $\mathcal{M}$ as candidates for good approximations to $\text{Buy}^k\text{Rev}(\mathcal{D})$ for the case $k \leq n^{1/2-\varepsilon}$.

4.1 Proof of Theorem 3

Throughout this Subsection we again abuse the fact that for additive valuation functions x , we can think of the valuation as an n -dimensional vector $\vec{x} = (x_1, \dots, x_n)$.

4.1.1 Part 1: Description of the Instance

In order to describe the instance we consider, we first need to introduce the concept of k -cover-free families of sets.

Definition 3. *A family of sets $\mathcal{F}$ is called k -cover-free if $A_0 \not\subseteq A_1 \cup A_2 \cup \dots \cup A_k$ holds for all distinct $A_0, A_1, \dots, A_k \in \mathcal{F}$.*

We will be interested in constructing the largest possible family of sets that is k -cover-free. Let $T(n, k)$ denote the maximum cardinality of a k -cover-free family of sets $\mathcal{F}$. We use the following bound from [16], attributed there to [21].

Theorem 19 ([21]). *For all n, k , it holds that*

$$\Omega\left(\frac{1}{k^2}\right) \leq \frac{\log(T(n, k))}{n}.$$

In other words, there exists a family of sets $\mathcal{F}_k$ that is k -cover-free and $|\mathcal{F}_k| \geq 2^{\Omega(\frac{n}{k^2})}$.

We will use k -cover-free sets to construct pairs of sequences that have large menu gaps. Then, we will take this pair of sequences and show how to obtain a n -dimensional distribution whose revenue gap is lower bounded by the menu gap of the underlying pair of sequences. We are now ready to define the instance of interest. Assume $k \leq n^{1/2-\varepsilon}$.

Definition 4. *Let $\mathcal{F}^L$ be a k -cover-free family of sets of maximal size, i.e., such that $|\mathcal{F}^L| = T(n, k) = \exp(\Omega(n/k^2))$. Set $\vec{x}_i^L = \vec{q}_i^L = \vec{e}_{A_i} \forall A_i \in \mathcal{F}^L$, where by $\vec{e}_S$ we denote the n -dimensional indicator vector for set S .*

Observe that unlike other constructions (e.g., [19], [26]) the number of points in each pair of sequences is finite. Thus this instance cannot witness an infinite revenue gap, but we claim it can witness an exponential revenue gap.

4.1.2 Part 2: The Instance has Large Menu Gap

In the next step of the proof of Theorem 3 we will show that the constructed instance has large menu gap.

Claim 20. *For the instance described in Definition 4, it holds that*

$$\text{MenuGap}^k(X^L, Q^L) \geq \frac{|\mathcal{F}^L|}{n}.$$

Proof.

$$\begin{aligned} \text{gap}_i^k(X^L, Q^L) &= \min_{j_1, j_2, \dots, j_k \leq i} \frac{\vec{e}_{A_i}}{|\vec{e}_{A_i}|} \cdot \left(\vec{e}_{A_i} - \text{Lot}(\vec{e}_{A_{j_1}}, \vec{e}_{A_{j_2}}, \dots, \vec{e}_{A_{j_k}}) \right) \\ &\geq \min_{j_1, j_2, \dots, j_k \neq i} \frac{\vec{e}_{A_i}}{|\vec{e}_{A_i}|} \cdot \left(\vec{e}_{A_i} - \text{Lot}(\vec{e}_{A_{j_1}}, \vec{e}_{A_{j_2}}, \dots, \vec{e}_{A_{j_k}}) \right) \\ &\geq \min_{j_1, j_2, \dots, j_k \neq i} \frac{\vec{e}_{A_i}}{|\vec{e}_{A_i}|} \cdot \left(\vec{e}_{A_i} - \vec{e}_{\cup_{\ell=1}^k A_{j_\ell}} \right) \\ &\geq \min_{j_1, j_2, \dots, j_k \neq i} \frac{|A_i| - |A_i \cap (\cup_{\ell=1}^k A_{j_\ell})|}{|A_i|} \\ &= \min_{j_1, j_2, \dots, j_k \neq i} \frac{|A_i \setminus (\cup_{\ell=1}^k A_{j_\ell})|}{|A_i|} \geq \frac{1}{n}. \end{aligned}$$

The first inequality observes that the gap only worsens when we allow for *all other* points to be used, rather than just those that come before i . The second inequality observes that, since all

vectors inside the argument have integral coordinates, the output is the indicator vector over the union of the inputs. The third inequality observes that for any two sets S, T , $\vec{e}_S \cdot \vec{e}_T = |S \cap T|$. The fourth inequality restates the previous line. The last inequality uses $|A_i| \leq n$ in the denominator and the fact that $\mathcal{F}^L$ is k -cover free, thus $A_i \setminus (\cup_{\ell=1}^k A_{j_\ell}) \neq \emptyset$ for any choice of A_{j_ℓ} in the numerator.

Thus, $\text{gap}_i^k(X^L, Q^L) \geq 1/n$ for all $i \in \mathcal{F}^L$. Therefore, $\text{MenuGap}^k(X^L, Q^L) \geq \frac{|\mathcal{F}^L|}{n}$. $\square$

4.1.3 Part 3: from Sequences to Distributions

We now present the final piece for the proof of Theorem 3. Lemma 21 states that given a pair of sequences (X, Q) of a certain form, we can find a distribution whose revenue gap is at least as large as $\text{MenuGap}^k(X, Q)$. This is a slight generalization of a lemma from cite [19]. The experienced reader will notice that our construction uses many similar ideas. Their work makes no assumptions on the sequences X, Q , but only works for the case of $k = 1$.

Lemma 21. *Let (X, Q) be a pair of sequences such that $\vec{x}_i \in \{0, 1\}^n, \vec{q}_i \in \{0, 1\}^n$ for all i . Moreover, suppose $\text{gap}_i^k(X, Q) \geq \frac{1}{n}$ for all i . Then, there exists a distribution $\mathcal{D} \in \mathbb{R}^n$ such that for any integer k ,*

$$\frac{\text{Buy}^k \text{Rev}(\mathcal{D})}{\text{BRev}(\mathcal{D})} \geq \frac{\text{MenuGap}^k(X, Q)}{2n}.$$

Proof. In order to construct a distribution $\mathcal{D}$ we need both a valuation $\hat{x}$ and a density function $f(\text{hat}x)$. Let $C_i = (n+1)^{2i}$. Then we define distribution $\mathcal{D}$ by setting $\vec{\hat{x}}_i = \vec{x}_i \cdot C_i, f(\vec{\hat{x}}_i) = \frac{1}{C_i}$. It is clear that this defines a valid distribution, i.e., $f(\vec{\hat{x}}_i) \geq 0$ and $\sum_i f(\vec{\hat{x}}_i) \leq 1$. Place the rest of the probability mass at a valuation of $\vec{0}^n$.

We will now show that $\text{Buy}^k \text{Rev}(\mathcal{D}) \geq \text{MenuGap}^k(X, Q)$ and $\text{BRev}(\mathcal{D}) \leq 2n$. Consider the menu $\mathcal{M}$ which offers allocation $\vec{q}_i$ at price $p_i = \text{gap}_i^k(X^L, Q^L) \cdot C_i$. Let $\mathcal{M}_i$ be the sub-menu of $\mathcal{M}$ consisting of the first i menu entries and the $(0, \vec{0}^n)$ entry. We will first claim that $\mathcal{M}$ is a buy- k menu.

Claim 22. *Any valuation $\hat{x}_i$ prefers to purchase the menu entry $(p_i, \vec{q}_i)$ to any other combination of k menu entries from $\mathcal{M}_i$.*

Proof. First, observe that if the valuation $\hat{x}_i$ purchases at least one copy of $(p_i, \vec{q}_i)$, because $\vec{x}_i = \vec{q}_i$ and $\vec{q}_i \in \{0, 1\}^n$, there is no value in purchasing any other menu entry. This is because a buyer with valuation $\hat{x}_i$ is only interested in the items in the support of $\vec{q}_i$, all of which are given to the buyer with probability 1. There is no benefit from purchasing any other lottery. Thus, any other reasonable deviations involve buying up to k menu entries from $\mathcal{M}_{i-1}$. The utility from purchasing any such combination is upper bounded by $\hat{x}_i \cdot \vec{\text{Lot}}(\vec{q}_{i_1}, \vec{q}_{i_2}, \dots, \vec{q}_{i_k})$. The utility from purchasing $(p_i, \vec{q}_i)$ is $\vec{\hat{x}}_i \cdot \vec{q}_i - p_i$. By choice of $p_i, \vec{\hat{x}}_i$ we get that this is

$$C_i \vec{x}_i \cdot \vec{q}_i - C_i \cdot \text{gap}_i^k(X^L, Q^L) \geq C_i \vec{x}_i \cdot \vec{\text{Lot}}(\vec{q}_{i_1}, \vec{q}_{i_2}, \dots, \vec{q}_{i_k}),$$

where the inequality follows from recalling the definition of $\text{gap}_i^k(X, Q)$ (and cancelling the C_i). $\square$

The next thing we need to show is that the valuation will not prefer to buy any option on $\mathcal{M} \setminus \mathcal{M}_i$. The utility from purchasing the preferred option is at most $C_i \cdot n$. The cost of any further option is at least $C_{i+1} \cdot \text{gap}_{i+1}^k(X^L, Q^L)$. By assumption, $\text{gap}_i^k(X, Q) \geq \frac{1}{n}$. Therefore, the price of

any option with $j > i$ is at least C_{i+1}/n . By construction, the price alone for any option $(p_j, \vec{q}_j)$ with $j > i$ is already greater than the possible utility the buyer could get. Thus, purchasing such menu entries would give them non-positive utility. Therefore, a valuation $\hat{x}_i$ will purchase exactly one copy of the menu entry $(p_i, \vec{q}_i)$. The revenue of mechanism $\mathcal{M}$ is $\sum_i f(\hat{x}_i)p_i = \sum_i \text{gap}_i^k(X, Q) = \text{MenuGap}^k(X, Q)$. Since $\mathcal{M}$ is a buy- k menu, $\text{Buy}^k\text{Rev}(\mathcal{D}) \geq \text{Buy}^k\text{Rev}(\mathcal{D}, \mathcal{M})$.

All that remains is to show that the revenue of bundling is at most $2n$. Note that the value a valuation $\hat{x}_i$ has for the bundle is at most $nC_i \leq C_{i+1}$. Thus, any price between $(C_{i-1} \cdot |\vec{x}_{i-1}|, C_i \cdot |\vec{x}_i|]$ will sell to the same set of bidders. Since we want to maximize revenue, it only makes sense to consider prices $b_i = C_i \cdot |\vec{x}_i|$ for all i . Consider any such price b_i for the bundle. The revenue is $b_i \cdot \Pr_{\hat{x}_j \sim \mathcal{D}}(C_j |\vec{x}_j| \geq b_i) = b_i \cdot \Pr_{\hat{x}_j \sim \mathcal{D}}(C_j \geq C_i) = b_i \cdot \sum_{j \geq i} f(\hat{x}_j) = b_i \cdot \frac{2n}{C_i(2n-1)} \leq 2b_i \cdot \frac{1}{C_i} \leq 2n$. $\square$

4.1.4 Part 4: Putting it all together

We are now ready to present the Proof of Theorem 3.

Proof of Theorem 3. Consider the instance (X^L, Q^L) from Definition 4. By Claim 20, the instance satisfies $\text{MenuGap}^k(X^L, Q^L) \geq |\mathcal{F}^L|/n$ and has only integral vectors. By Lemma 21, we can turn the pair of sequences into a distribution $\mathcal{D}$ with $\text{Buy}^k\text{Rev}(\mathcal{D})/\text{BRev}(\mathcal{D}) \geq \frac{\text{MenuGap}^k(X^L, Q^L)}{2n}$. Thus,

$$\frac{\text{Buy}^k\text{Rev}(\mathcal{D})}{\text{BRev}(\mathcal{D})} \geq \frac{|\mathcal{F}^L|}{2n^2}.$$

Finally, for $k \leq n^{1/2-\varepsilon}$, note that $|\mathcal{F}^L| = \exp(\Omega(n^\varepsilon))$. Therefore, the right hand side becomes $\frac{\exp(\Omega(n^\varepsilon))}{2n^2}$. $\square$

5 Conclusion

In this paper we initiate the study of fine-grained buy-many mechanisms. The motivation for our work stems from a simple observation: there exist distributions for which the buy-one revenue gap $\text{Rev}(\mathcal{D})/\text{BRev}(\mathcal{D})$ is unbounded, but for all distributions the buy-many revenue gap $\text{BuyManyRev}(\mathcal{D})/\text{BRev}(\mathcal{D})$ is finite. There is a wide worst-case revenue gap between optimal buy-many and optimal buy-one mechanisms, which begs the question: how much must we constraint the seller's choice of mechanism until the revenue gap becomes finite for all distributions? In order to answer this question, we introduce the concept of buy- k mechanisms, those where the buyer can buy any multi-set of up to k many menu choices. We show that buy- n mechanisms are not much better than bundling. For all distributions $\mathcal{D}$, the revenue from bundling recovers a $O(n^2)$ fraction of the optimal buy- n revenue when the buyer is additive and a $O(2^n \cdot n^2)$ fraction of the optimal buy- n revenue when the buyer has an arbitrary monotone valuation. Our proof uses a recent framework proposed in [19, 26] for buy-one mechanisms. While in those works, the framework has been used to produce examples of inapproximable distributions, our work shows that it can be used to prove approximation guarantees. Moreover, all our results hold for the case of an adaptive buyer.

There are numerous questions for future work:

- We have already outlined one interesting question for future work, to prove or disprove Conjecture C. Even a slight weakening of the conjecture, just showing that for every k there exists some $\mathcal{D}_k$ such that $\text{Buy}^k\text{Rev}(\mathcal{D}_k) > \text{Buy}^{k+1}\text{Rev}(\mathcal{D}_k)$ would be interesting. We discuss our candidate instance for Conjecture 1 in Appendix C.

- We showed that restricting the seller to be buy- n incentive compatible sufficed to obtain a $O(2^n \cdot n^2)$ -approximation via bundling. It would be interesting if the exponential bound is tight in general; if it is, it would be interesting to characterize the distributions for which a polynomial approximation is possible.
- It would be interesting to understand the role of k in whether or not the revenue gap is finite. Concretely, we would like to answer the following question: for a given n , what is the smallest k for which $\text{Buy}^k\text{Rev}(\mathcal{D})/\text{BRev}(\mathcal{D})$ is finite for all $\mathcal{D}$? Ultimately, we would like to understand the exact trade off between k, n in the revenue gap, answering Open Question 1.
- Future work could also follow the steps of [14] in understanding whether or not fine-grained buy-many mechanisms satisfy revenue monotonicity, or whether or not fine-grained buy-many mechanisms admit $(1 - \varepsilon)$ -approximations via finite-sized mechanisms (and what role k has in answering any of these questions).
- Another interesting avenue would be to explore the power that buy-many or fine-grained buy-many mechanisms have over product distributions. There is a long line of work with elegant approximation results for the case of product distributions, but progress towards polynomial time approximation schemes has been slow. It is possible that restricting the seller's choice of mechanism improves the performance of existing algorithms or allows for the discovery of more efficient ones.
- Finally, very little is understood computationally about buy-many and fine-grained buy-many mechanisms. For instance, it is not immediately clear how to efficiently test whether or not a mechanism is buy-many or buy- k for some k .

We hope that our results strengthen the importance of developing a deeper understanding of fine-grained buy-many mechanisms.

References

- [1] Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. 2017. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 869–877. <https://doi.org/10.1145/3055399.3055426>
- [2] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- [3] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2010. Pricing Randomized Allocations. In *the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [4] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2015. Pricing lotteries. *J. Economic Theory* 156 (2015), 144–174. <https://doi.org/10.1016/j.jet.2014.04.011>

- [5] Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. 2016. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the 48th ACM Conference on Theory of Computation (STOC)*.
- [6] Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 170–183. <https://doi.org/10.1145/3055399.3055465>
- [7] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*.
- [8] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.
- [9] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. 2015. The power of randomness in Bayesian optimal mechanism design. *Games and Economic Behavior* 91 (2015), 297–317. <https://doi.org/10.1016/j.geb.2012.08.010>
- [10] Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*. 579–596. <https://doi.org/10.1145/2940716.2940756>
- [11] Shuchi Chawla, Rojin Rezvan, Yifeng Teng, and Christos Tzamos. 2022. Pricing ordered items. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, Stefano Leonardi and Anupam Gupta (Eds.). ACM, 722–735. <https://doi.org/10.1145/3519935.3520065>
- [12] Shuchi Chawla, Yifeng Teng, and Christos Tzamos. 2019. Buy-Many Mechanisms Are Not Much Better than Item Pricing. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC '19)*. Association for Computing Machinery, New York, NY, USA, 237–238. <https://doi.org/10.1145/3328526.3329583>
- [13] Shuchi Chawla, Yifeng Teng, and Christos Tzamos. 2020. Buy-Many Mechanisms: What Are They and Why Should You Care? *SIGecom Exch.* 18, 1 (dec 2020), 12–18. <https://doi.org/10.1145/3440959.3440963>
- [14] Shuchi Chawla, Yifeng Teng, and Christos Tzamos. 2020. Menu-Size Complexity and Revenue Continuity of Buy-Many Mechanisms. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC '20)*. Association for Computing Machinery, New York, NY, USA, 475–476. <https://doi.org/10.1145/3391403.3399453>
- [15] Constantinos Daskalakis. 2015. Multi-Item Auctions Defying Intuition? *SIGecom Exch.* 14, 1 (nov 2015), 41–75. <https://doi.org/10.1145/2845926.2845928>
- [16] Zoltán Füredi. 1996. Onr-Cover-Free Families. *J. Comb. Theory Ser. A* 73, 1 (jan 1996), 172–173. <https://doi.org/10.1006/jcta.1996.0012>

- [17] Sergiu Hart and Noam Nisan. 2012. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*.
- [18] Sergiu Hart and Noam Nisan. 2013. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*.
- [19] Sergiu Hart and Noam Nisan. 2017. Approximate revenue maximization with multiple items. *J. Economic Theory* 172 (2017), 313–347. <https://doi.org/10.1016/j.jet.2017.09.001>
- [20] Sergiu Hart and Philip J. Reny. 2015. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics* 10, 3 (2015), 893–922.
- [21] W. Kautz and R. Singleton. 1964. Nonrandom binary superimposed codes. *IEEE Transactions on Information Theory* 10, 4 (1964), 363–377. <https://doi.org/10.1109/TIT.1964.1053689>
- [22] Pravesh Kothari, Sahil Singla, Divyarthi Mohan, Ariel Schwartzman, and S. Matthew Weinberg. 2019. Approximation Schemes for a Unit-Demand Buyer with Independent Items via Symmetries. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, David Zuckerman (Ed.). IEEE Computer Society, 220–232. <https://doi.org/10.1109/FOCS.2019.00023>
- [23] Xinye Li and Andrew Chi-Chih Yao. 2013. On Revenue Maximization for Selling Multiple Independently Distributed Items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.
- [24] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [25] Alexandros Psomas, Ariel Schwartzman, and S. Matthew Weinberg. 2019. Smoothed Analysis of Multi-Item Auctions with Correlated Values. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC ’19)*. Association for Computing Machinery, New York, NY, USA, 417–418. <https://doi.org/10.1145/3328526.3329563>
- [26] Alexandros Psomas, Ariel Schwartzman, and S. Matthew Weinberg. 2022. On Infinite Separations Between Simple and Optimal Mechanisms: A Converse of a Theorem of Hart and Nisan. *Manuscript* (2022).
- [27] Aviad Rubinstein and S. Matthew Weinberg. 2015. Simple Mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the 16th ACM Conference on Electronic Commerce*.
- [28] Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

A Adaptive Buy-Many Mechanisms

In this Appendix, we briefly review another notion of buy- k mechanisms, which we refer to as *adaptive* buy- k mechanisms. We will define them to be analogues of the adaptive buy-many mechanisms as defined in [12].

In the standard definition of buy- k mechanisms, formalized in section 2, the buyer may purchase any multi-set of menu options of size up to k . In a randomized mechanism, this corresponds to committing to up to k options and *only after that* receiving their outcome allocations; in other words, the *choice* of, say, second option, is *not* a function of probabilistic outcomes of the lottery for the first option. This is formally captured in our definition of the function $\vec{\text{Lot}}(\cdot)$.

We can naturally also consider a variant of this definition that allows for *adaptively* choosing the options to purchase, based on the probabilistic outcomes of the lotteries for the prior options. In this case, the buyer can commit to a *strategy* of different ways of purchasing up to k options, while seeing the outcome of each purchased lottery before purchasing the next option. A strategy can be thought of as a 2^n -ary tree of depth at most k where each node identifies what to purchase on the next step depending on which items of the current purchased lottery “succeeded” or “failed”. The buyer is then interested in a strategy with maximum expected payoff. Analogous to [12], we say a mechanism $\mathcal{M}$ is *adaptive buy- k incentive-compatible* if for every valuation of the buyer, the strategy with maximum expected payoff consists of buying a single option (see also Section 2 of [12] for more details on this definition).

As was observed in [12], it is easy to see that since the set of non-adaptive buy- k options are all valid strategies for an adaptive buy- k mechanism, any mechanism that is adaptive buy- k incentive-compatible is also (non-adaptive) buy- k incentive-compatible (but the reverse direction is not necessarily true). As a corollary of this, we can immediately extend our bounds in Theorems 1 and 2 to adaptive buy- k incentive-compatible mechanisms.

Finally recall that the construction of Theorem 3 presented in Section 4 gave a deterministic mechanism. When a mechanism is deterministic, there is no distinction between adaptive and non-adaptive strategies because there is no randomness in the allocation. Therefore, the lower bounds of Theorem 3 also extend to the case of adaptive buyers.

B $\text{BRev}(\mathcal{D})$ Can Not Give a Sublinear Approximation to $\text{Buy}^k\text{Rev}(\mathcal{D})$

In this section we show that Theorem 1 can not be improved to a sublinear approximation factor for any k .

Claim 23. *There exists a distribution $\mathcal{D}$ such that for a single, additive buyer and any k ,*

$$\text{BRev}(\mathcal{D}) \leq \frac{2 \cdot \text{Buy}^k\text{Rev}(\mathcal{D})}{n}.$$

Proof. Consider the distribution $\mathcal{D}$ from Example 15 of [17]. This distribution satisfies the following property: $\text{BRev}(\mathcal{D}) = 2, \text{SRev}(\mathcal{D}) = n$, where $\text{SRev}(\cdot)$ is the optimal revenue attained by item-pricing. Observe that item-pricing is a buy-many incentive-compatible mechanism. Thus, $\text{SRev}(\mathcal{D}) \leq \text{BuyManyRev}(\mathcal{D})$. Moreover, from Claim 1 we know that for any k , $\text{BuyManyRev}(\mathcal{D}) \leq \text{Buy}^k\text{Rev}(\mathcal{D})$. Therefore,

$$\text{BRev}(\mathcal{D}) \leq \frac{2 \cdot \text{Buy}^k\text{Rev}(\mathcal{D})}{n}.$$

In particular, this shows that the ratio between $\text{BRev}(\mathcal{D})$ and $\text{Buy}^k\text{Rev}(\mathcal{D})$ for additive buyers is $\Omega(n)$, implying that Theorem 1 can not be substantially improved. $\square$

C Candidate Distribution for a Separation Between $\text{Buy}^k\text{Rev}(\mathcal{D})$ and $\text{Buy}^{k+1}\text{Rev}(\mathcal{D})$

In this section we present the instance $\mathcal{D}$ that proves Proposition 1 and that we posit could prove Conjecture 1.

Example 2. Consider the following (correlated) distribution $\mathcal{D}$ over two additive items:

$$\Pr(v_1 = a, v_2 = b) = \begin{cases} 1/6 & \text{for } a = 3, b = 4 \\ 1/6 & \text{for } a = 4, b = 3 \\ 4/6 & \text{for } a = 5, b = 7. \end{cases}$$

For the following proofs, we introduce some additional notation. Namely, we denote each $x_j \in \text{SUPP}(\mathcal{D})$ as “buyer j ” and let $u_j(\Lambda) = x_j(\vec{\text{Lot}}(\Lambda)) - \sum_{i=1} p(\vec{q}_i)$ be the utility gained by buyer j from purchasing the multi-set of allocations Λ . We also prove the following useful claim about the utility function.

Claim 24. Let buyer j have a valuation of x_j . Let Λ be some multi-set of allocations such that $q_i \in \Lambda$ with multiplicity $m_i \geq 0$. Then, $u_j(\Lambda)$ is concave in m_i .

Proof. We verify that the second derivative of the utility function for buyer j with respect to m_i is non-positive.

$$\begin{aligned} \frac{\partial^2 u_j(\Lambda)}{\partial m_i^2} &= \frac{\partial^2}{\partial m_i^2} \left(x_j(\vec{\text{Lot}}(\Lambda)) - \sum_{i=1} p(\vec{q}_i) \right) \\ &= \frac{\partial^2}{\partial m_i^2} \left(\sum_{t=1} x_{jt} \cdot (1 - (1 - q_{1t})^{m_1} \cdot \dots \cdot (1 - q_{it})^{m_i}) - \sum_{i=1} p(\vec{q}_i) \right) \\ &= \sum_{t=1} -x_{jt} \cdot \ln(1 - q_{it})^2 \cdot (1 - q_{1t})^{m_1} \cdot \dots \cdot (1 - q_{it})^{m_i} \\ &\leq 0 \end{aligned}$$

since $\vec{q}_i \in [0, 1]$ and $x_j \geq 0$. $\square$

Proof of Proposition 1. The LP formulation presented in [4] provides us with a simple way to compute optimal mechanisms in the buy-one world. Naturally, for some fixed input distribution, the LP constructs an allocation and payment function which maximizes expected revenue while maintaining feasibility and buy-one incentive-compatibility constraints. Through this LP, we can compute the revenue-optimal buy-one mechanism for the distribution $\mathcal{D}$ shown below:

$$\mathcal{M}_1 = \begin{cases} ((0.2, 0.2), 1.4) \\ ((1, 0), 4) \\ ((1, 1), 11). \end{cases}$$

This mechanism achieves an expected revenue (subject to truncation) of 8.233. More generally, the previous program can be altered to compute the optimal buy- k mechanism for a distribution $\mathcal{D}$, albeit by introducing non-convex constraints to enforce buy- k incentive-compatibility. Despite non-convexity, for $k = 2, 3, 4$, a non-convex optimizer was able to compute the revenue-optimal buy- k mechanism. The annotated code samples for computing these optimal mechanisms can be found here. For each respective value of k , the optimizer found the optimal expected revenue to be $[8.135, 8.096, 8.074]$, allowing us to conclude that $\text{Buy}^1\text{Rev}(\mathcal{D}) > \text{Buy}^2\text{Rev}(\mathcal{D}) > \text{Buy}^3\text{Rev}(\mathcal{D}) > \text{Buy}^4\text{Rev}(\mathcal{D})$. Thus, we empirically validate Proposition 1. $\square$

We also conjecture that Example 2 is a good candidate for proving Conjecture 1. Fix some value of $k \geq 1$ and consider the following mechanism $\mathcal{M}_k = \{t_1 = ((\alpha_k, \alpha_k), 7\alpha_k), t_2 = ((1, 0), 4), t_3 = ((1, 1), 11)\}$, where α_k is the smallest positive real root of the k -degree polynomial $f_k(x) = 12(1 - (1 - x)^k) - 7k \cdot x - 1$. Intuitively, the value α_k has the following property: when buyer 3 purchases k copies of t_1 , they receive a utility of 1, which is the same utility gained from purchasing a single copy of t_3 .

Claim 25. *The mechanism $\mathcal{M}_k$ is buy- k incentive compatible for the distribution defined in Example 2, and achieves expected revenue $R_k = 8 + 7/6 \cdot \alpha_k$.*

Proof. We first verify that $\mathcal{M}_k$ is a buy- k incentive-compatible mechanism for the distribution $\mathcal{D}$. By Claim 24 the utility of buyer j from purchasing a multi-set of allocations Λ is concave in the number of identical allocations bought. Consequently, once we establish that a buyer achieves non-positive utility from purchasing a multi-set Λ of allocations, we can conclude that purchasing any multi-set $\Lambda' \supseteq \Lambda$ will similarly yield non-positive utility for the buyer.

We proceed by calculating the utility received by each buyer for each allocation and showing that each buyer (weakly) maximizes their utility by purchasing a single ticket. First, consider buyer 1 who achieves $u_1((\alpha_k, \alpha_k)) = 0$, $u_1((1, 0)) < 0$, and $u_1((1, 1)) < 0$. Second, consider buyer 2 who achieves $u_2((\alpha_k, \alpha_k)) = 0$, $u_2((1, 0)) = 0$, and $u_2((1, 1)) < 0$. By Claim 24, both buyer 1 and buyer 2 satisfy the buy- k incentive-compatibility constraints since they will always prefer to purchase tickets t_1 and t_2 , respectively, compared to another other multi-set of allocations.

For buyer 3, a slightly more careful analysis is required. We can first check that buyer 3 satisfies buy-one incentive-compatibility constraints since $0 < u_3((\alpha_k, \alpha_k)) \leq 1$, $u_3((1, 0)) = 1$, and $u_3((1, 1)) = 1$. We must additionally check that the higher order incentive-compatibility constraints also hold since buyer 3 obtains positive utility from individually buying t_1 and t_2 . Notice, t_2 offers a deterministic allocation, so purchasing multiple copies of this ticket does not yield additional utility. This leaves two cases to analyze. Specifically, we can easily verify that $u_3(\{(\alpha_k, \alpha_k), (1, 0)\}) = 1$ and $u_3(\{(\alpha_k, \alpha_k)_1, \dots, (\alpha_k, \alpha_k)_k\}) = 1$, where the last equality follows directly from the definition of α_k . Moreover, since α_k is the smallest positive root of the polynomial $f_k(x)$, we know that $u_3(\{(\alpha_k, \alpha_k)_1, \dots, (\alpha_k, \alpha_k)_i\}) < 1$ for $1 \leq i < k$. By Claim 24, we can conclude that buyer 3 can achieve a utility of at most 1 by deviating from the ticket t_3 . Thus, buyer 3 satisfies buy- k incentive-compatibility constraints.

The revenue of $\mathcal{M}_k$ follows from the fact that buyer j will purchase ticket t_j and from the density of the valuation classes. In the case where two allocations yield the same utility for a buyer, we break the tie in favor of the seller. As a result, $R_k = 2/3 \cdot 11 + 1/6 \cdot 4 + 1/6 \cdot 7\alpha_k = 8 + 7/6 \cdot \alpha_k$. $\square$

Claim 26. *For all integers $k \geq 1$, the polynomial $f_k(x)$ always has a positive, real root below 1. Moreover, the sequence $\{\alpha_k\}_{k=1}^\infty$, where α_k is the smallest positive, real root of $f_k(x)$, is strictly decreasing.*

Proof. The proof follows via an inductive argument using the intermediate value theorem. First, recognize that f is continuous and $f_k(0) = -1$ for all $k \geq 1$. By the proof of Proposition 1, we have that $f_1(0.2) = 0 < 1$. Let us inductively assume that $f_k(\alpha_k) = 0$ and $\alpha_k \leq 1$ for $k \geq 1$. We wish to show that $f_{k+1}(\alpha_k) > 0$ as this would imply that $f_{k+1}(x)$ has a root $0 < \alpha_{k+1} < \alpha_k$ by the intermediate value theorem. To begin,

$$\begin{aligned} f_{k+1}(\alpha_k) &= -12(1 - (1 - \alpha_k)^{k+1}) - 7\alpha_k \cdot (k+1) - 1 \\ &= -12(1 - \alpha_k)^{k+1} + 12(1 - \alpha_k)^k - 7\alpha_k \\ &= -12(1 - \alpha_k)^k \alpha_k - 7\alpha_k, \end{aligned}$$

where the second equality follows from the fact that α_k is a root of $f_k(x)$. From the reduced form above, it suffices to show that $\alpha_k < (1 - (\frac{7}{12})^{1/k})$ to prove that $f_{k+1}(\alpha_k) > 0$. This can be accomplished by the following algebraic manipulations:

$$\begin{aligned} f_k(1 - (7/12)^{1/k}) &= -12((7/12)^{1/k})^k - 7k \cdot (1 - (7/12)^{1/k}) + 11 \\ &= -7k \cdot (1 - (7/12)^{1/k}) + 4 \\ &> 0. \end{aligned}$$

Since $f_k(1 - (7/12)^{1/k}) > 0$, by the intermediate value theorem, it follows that $\alpha_k < (1 - (\frac{7}{12})^{1/k})$. Consequently, we find that $f_{k+1}(\alpha_k) > 0$, finishing the proof that $0 < \alpha_{k+1} < \alpha_k < 1$. $\square$

Note that Claim 26 clearly implies that the sequence $\{R_k\}_{k=1}^{\infty}$ is also strictly decreasing as it is the same sequence as $\{c_k\}_{k=1}^{\infty}$ but shifted by a constant. Given Claims 25, 26, all that remains is to show that $\text{Buy}^k \text{Rev}(\mathcal{D}) = R_k$ for all $k \geq 2$. Then since $\text{Buy}^1 \text{Rev}(\mathcal{D}) > R_2$ and the sequence R_k is strictly decreasing, we would have that $\text{Buy}^k \text{Rev}(\mathcal{D}) > \text{Buy}^{k+1} \text{Rev}(\mathcal{D})$ for all k , proving Conjecture 1. The principal obstacle to proving the missing step is find a technique to overcome non-convexity of the optimization problem. We currently have a candidate mechanism, but the non-convexity of the constraints hinders our ability to prove it is optimal via some notion of duality.