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

Sepehr Assadi
sepehr.assadi@rutgers.edu

Ariel Schvartzman
as3569@dimacs.rutgers.edu

February 10, 2022

Abstract

Multi-item 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, even for the case of just two items and a single buyer [4, 18]. One possible explanation for this bizarre property is that the seller is unrestricted in their choice of mechanisms.

We introduce a new class of mechanisms, buy-k mechanisms, which smoothly interpolates between the classical buy-one mechanisms and buy-many mechanisms [11, 13, 12]. Buy-k mechanisms allow the buyer to (non-adaptively) buy up to k many menu options, progressively shrinking the seller's feasible set of mechanisms. We show that restricting the seller to the class of buy-n mechanisms suffices to overcome the bizarre, infinite revenue properties of the buy-one model for the case of a single, additive buyer. The revenue gap with respect to bundling, an extremely simple mechanism, is bounded by $O(n^3)$ for any arbitrarily correlated distribution $\mathcal{D}$ over n items. For the special case of $n = 2$, we show that the revenue-optimal buy-2 mechanism gets no better than 40 times the revenue from bundling. Our upper bounds also hold for the case of adaptive buyers due to an observation from [11].

Finally, we show that allowing the buyer to purchase a small number of menu options *does not suffice* to guarantee sub-exponential approximations. If the buyer is only allowed to buy $k = \Theta(n^{1/2-\varepsilon})$ many menu options, the gap between the revenue-optimal buy-k mechanism and bundling may be exponential in n. This implies that no "simple" mechanism can get a sub-exponential approximation in this regime. Moreover, our lower bound instance, based on combinatorial designs and cover-free sets, uses a buy-k deterministic mechanism. This allows us to extend our lower bound to the case of adaptive buyers.

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 private values, seminal work of Myerson [23] 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.

Multi-item optimal auctions are known to be complex objects, offering no discernible mathematical structure and often exhibiting "intuition defying" properties [14]. 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 [17, 18].[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.

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 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 need not satisfy this property, making them less appealing for real-world applications.

Consider the following concrete example. 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 work of [3, 4] already exhibits how buy-many mechanisms overcome the revenue gap problem: the authors 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.

[1] By "simple" mechanisms, we mean mechanisms of size polynomial in the number of items.

This was later extended to arbitrary valuations by [11], while preserving the approximation factor. The key to these results is not that item-pricing is a particularly good mechanism, but 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.

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 and thus can be approximated via "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 $\mathrm{BRev}(\cdot)$ (henceforth bundling) is often a benchmark of interest.[2]

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.[3] 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 $\mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D}) = \max_{\mathcal{M}\in\mathcal{B}_k(\mathcal{D})} \mathrm{Rev}(\mathcal{D},\mathcal{M})$ be the optimal revenue attainable by a buy-k incentive-compatible mechanism. Formally, we want to answer the following question.

Open Question 1. *Given integers n, k, when does* $\textsc{poly}(n,k)\cdot\mathrm{BRev}(\mathcal{D}) \geq \mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D})$ *hold for all distributions $\mathcal{D}$ over n items?*

1.1 Our Contributions

We answer Open Question 1 affirmatively for the case when $k = n$ and show that $O(n^3)$ suffices. Our first main result shows that restricting the seller to the class of buy-two mechanisms suffices to get around pathological constructions for two items (like e.g., [18, 25]). These works show that there are distributions over 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 exists (even if we allowed super-polynomial ones). We show that the revenue from optimally pricing the bundle of the two items, $\mathrm{BRev}(\mathcal{D})$, recovers a constant-factor of the optimal buy-two revenue.

[2]Bundling is arguably one of the simplest mechanisms.

[3]In 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 B for a more detailed discussion.

Theorem 2. *For any distribution $\mathcal{D}$ over 2 items, it holds that*

$$40 \cdot \mathrm{BRev}(\mathcal{D}) \geq \mathrm{Buy2Rev}(\mathcal{D}).$$

Our next main result extends Theorem 2 to the case of n items.

Theorem 3. *For any distribution $\mathcal{D}$ over n items, it holds that*

$$O(n^3) \cdot \mathrm{BRev}(\mathcal{D}) \geq \mathrm{Buy}(n)\mathrm{Rev}(\mathcal{D}).$$

Taken together, Theorems 2, 3 show that $O(n^3)$ suffices for the case when $n = k$, partially answering Open Question 1. There are two subtle implications of these results. The first is that for all n-dimensional distributions $\mathcal{D}$, $\mathrm{Buy}(n)\mathrm{Rev}(\mathcal{D})$ is *finite* whenever $\mathrm{BRev}(\mathcal{D})$ is finite. This stands in contrast to the buy-one case where even for just $n = 2$ items, there exist $\mathcal{D}$ such that $\mathrm{Buy}(1)\mathrm{Rev}(\mathcal{D}) > \infty$ but $\mathrm{BRev}(\mathcal{D}) = O(1)$. The second is that since $\mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D}) \geq \mathrm{Buy}(k')\mathrm{Rev}(\mathcal{D})$ whenever $k < k'$,[4] then Theorems 2, 3 in fact also cover the case when $n \leq k$.

The proofs of Theorems 2, 3 will follow exactly the same outline. We present them separately for ease of exposition. Key to both theorems is the identification of a measure, $\mathrm{MenuGap}^k(\cdot,\cdot)$, whose formal definition we defer to Section 2. This quantity is the generalization to buy-k mechanisms of $\mathrm{MenuGap}(\cdot,\cdot)$ introduced by previous work for buy-one mechanisms (see [18, 25]). In those works, $\mathrm{MenuGap}(\cdot,\cdot)$ was used to construct distributions whose optimal revenue was hard to approximate. In contrast, our work shows that their framework can be used to prove approximation guarantees.

The first piece of the proof for Theorems 2 (resp. Theorem 3) is to show that there exists an appropriate choice of inputs (X, Q) such that $\mathrm{MenuGap}^2(X, Q)$ (resp. $\mathrm{MenuGap}^k(X, Q)$) upper bounds the ratio between the optimal buy-two (resp. buy-k) revenue and the revenue achieved by bundling, up to some constant factor (resp. up to some $O(k^2)$ factor). The second step of the proof upper bounds $\mathrm{MenuGap}^2(X, Q)$ itself by 2 (resp. $\mathrm{MenuGap}^n(X, Q)$ by n for the case $k = n$) for *all* input pairs (X, Q).

The next goal is to answer Open Question 1 for the case when $1 < k < n$. We make progress by giving a negative answer 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 4.

Theorem 4. *If $k \leq n^{1/2-\varepsilon}$ for some $\varepsilon > 0$, then there exists a distribution $\mathcal{D}$ over n items such that*

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

The careful reader will observe that Theorem 4 says something strong about the the revenue guarantees that "simple" mechanisms can obtain. Due to a

[4] See Claim 5 in Section 2 for a formal proof.

Corollary from [18] 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 $\textsc{poly}(n)$ cannot exceed $\textsc{poly}(n) \cdot \mathrm{BRev}(\mathcal{D})$. Thus, Theorem 4 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 4 will, unsurprisingly, borrow ideas from [18]. 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.

A Note on Adaptive vs Non-Adaptive Buyers. 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 2, 3) and our lower bounds (Theorem 4) to the case of adaptive buyers. We defer this discussion to Appendix B.

Our Techniques. 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 [18, 25] to the buy-k setting. Similar to [25], we prove that this measure upper bounds the revenue gap between the revenue-optimal mechanism (in some class of mechanisms) and bundling. Unlike [25], we are able to show a *finite* upper bound for this measure in the case of buy-n mechanisms, yielding a finite approximation result. We believe these ideas can be further used for other settings and valuation classes beyond additive.

1.2 Related Work

Buy-many mechanisms have been proposed as early as [3, 4]. Results from a recent line of work [11, 13, 12, 10] make the case to further the study of buy-many mechanisms. For instance, [13] 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 [19, 24]. In addition, as mentioned earlier in the introduction, [11] show that item-pricing recovers a $O(\log n)$ factor of the optimal buy-many revenue. Combining a Corollary from [18] with the main result of [11] 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 Claim 5). 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.

[11] 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, [13] 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., [22, 16, 27, 6, 8, 2, 5, 26, 7, 9, 1], among others), effectively circumventing the pathological constructions of [18]. 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, [21] 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 [24] provides yet another way to circumvent the pathological constructions of [18]. 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 proof of Theorem 2. Section 4 contains the proof of Theorem 3. Section 5 contains the proof of Theorem 4. We conclude in section 6 and outline questions for future work. Appendix A includes some omitted proofs, while Appendix B contains the discussion for the case of adaptive buyers.

2 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 vector $\vec{v} = (v_1, \ldots, v_n)$ from a known, possibly correlated n-dimensional distribution $\mathcal{D}$, with support set $\mathcal{T} = \textsc{supp}(\mathcal{D})$. We assume that the buyer is additive across the items, i.e., for any subset of items $S \subseteq [n]$, $\vec{v}(S) = \sum_{i \in S} v_i$, and monotone, i.e., whenever $S \subseteq T$, $\vec{v}(S) \leq \vec{v}(T)$. Given a possibly randomized allocation of items $\vec{q} \in [0,1]^n$, we use $\vec{v}(\vec{q})$ to denote the buyer's expected utility: $\vec{v}(\vec{q}) = \sum_{i=1}^n v_i \cdot q_i$. Let $\Lambda = \{\vec{q}_1, \vec{q}_2, \ldots, \vec{q}_k, \ldots\}$ be a multi-set of allocations (of possibly unbounded size). We denote by $\vec{\mathrm{Lot}}(\Lambda) \in [0,1]^n$ (read "lottery") the expected allocation that results from being allocated every $\vec{q}_i \in \Lambda$ independently and at once, i.e., for all $j \in [n]$, $\vec{\mathrm{Lot}}(\Lambda)_j = 1 - \prod_{i=1}(1 - q_{ij})$.[5]

A mechanism $\mathcal{M} = (p, q)$ is defined by a pair of functions $p : \mathcal{T} \to \mathcal{R}_{\geq 0}$, $q : \mathcal{T} \to [0,1]^n$ known as the pricing and allocation functions, respectively. For

[5]The 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 B, but the main takeaway is that our upper bounds also hold for adaptive buyers.

a fixed integer k we say that a mechanism $\mathcal{M}$ is *buy-k* incentive compatible if for every valuation $\vec{v} \in \mathcal{T}$ 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, if $\vec{v}(\vec{q}(\vec{v}))-p(\vec{v}) \geq \vec{v}(\vec{\mathrm{Lot}}(\Lambda))-\sum_{i\in\Lambda} p(\vec{q}_i)$ 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 $k \to \infty$ 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 distribution $\mathcal{D}$, let

- $\mathrm{BRev}(\mathcal{D})$ be the revenue of the mechanism which sells the grand bundle for its optimal price. Namely, $\mathrm{BRev}(\mathcal{D}) = \max_p p \cdot \Pr_{\vec{v}\sim\mathcal{D}}(\sum v_i \geq p)$,
- $\mathrm{Rev}(\mathcal{D})$ be the revenue of the optimal buy-one incentive-compatible mechanism,
- $\mathrm{Rev}(\mathcal{D}, \mathcal{M})$ be the revenue of mechanism $\mathcal{M}$ when the buyer is allowed to buy up to 1 menu entry from $\mathcal{M}$.
- $\mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D})$ be the revenue of the optimal buy-k incentive-compatible mechanism,
- $\mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D}, \mathcal{M})$ 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 $\mathcal{M}$.
- $\mathrm{BuyManyRev}(\mathcal{D})$ be the revenue of the optimal buy-many incentive compatible mechanism.

Claim 5.

$$\mathrm{Rev}(\mathcal{D}) = \mathrm{Buy}(1)\mathrm{Rev}(\mathcal{D}) \geq \mathrm{Buy}(2)\mathrm{Rev}(\mathcal{D}) \geq \cdots \geq \mathrm{BuyManyRev}(\mathcal{D}) \geq \mathrm{BRev}(\mathcal{D}).$$

Proof. 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. □

2.2 Menu Gaps: An Intermediary Measure

We now present the definition of $\mathrm{gap}_i^k(X, Q)$ and $\mathrm{MenuGap}^k(X, Q)$, the quantities that will serve as intermediaries in proving Theorems 2, 3.

Definition 6. *Let* $X = \{\vec{x}_i\}_{i=1}^N \in \mathbb{R}^k_{\geq 0}, Q = \{\vec{q}_i\}_{i=0}^N \in [0,1]^k$ *be sequences of vectors with* $\vec{q}_0 = \vec{0}^k$. *Then*

$$\text{gap}_i^k(X, Q) = \min_{j_1, j_2, \ldots, j_k < i} \vec{x}_i \cdot (\vec{q}_i - \vec{\text{Lot}}(\vec{q}_{j_1}, \vec{q}_{j_2}, \ldots, \vec{q}_{j_k})), \tag{1}$$

and

$$\text{MenuGap}^k(X, Q) = \sum_{i=1}^N \text{gap}_i^k(X, Q)/||\vec{x}_i||_1. \tag{2}$$

These measures are generalizations of similar notions introduced in [18] and further developed by [25]. For the case where $k = 1$, we exactly recover these earlier definitions. In Definition 6, 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 $p_i = \text{gap}_k^i(X, Q)$ as the largest price a seller can post on menu entry $(p_i, \vec{q}_i)$ so that a buyer with valuation $\vec{v}_i$ 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 $\Lambda = \{\vec{q}_1, \vec{q}_2, \ldots, \vec{q}_k\}$, then $\vec{\text{Lot}}(\Lambda)$ dominates the vector which captures the coordinate-wise max entries of the vectors in Λ.

Claim 7. *If* $\Lambda = \{\vec{q}_1, \vec{q}_2, \ldots, \vec{q}_k\}, \vec{q}_i \in [0,1]^n$ *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 $(1 - q_{1j}) \cdot (1 - \prod_{j=2}^k (1 - q_{ij})) \geq 0$ since each term on the left is non-negative. Expanding it we get that $1 - q_{1j} - \prod_{j=1}^k (1 - q_{ij}) \geq 0$. 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 8. *Let* X, Q *be sequences as defined in Definition 6, 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 $(\vec{x}_i, \vec{q}_i)$ 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 $(\vec{x}_j, \vec{q}_j)$, 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.

□

3 Warm-up: Proof of Theorem 2

In this section we will prove Theorem 2, restated here for convenience.

Theorem 2. *For any distribution $\mathcal{D}$ over 2 items, it holds that*

$$40 \cdot \mathrm{BRev}(\mathcal{D}) \geq \mathrm{Buy2Rev}(\mathcal{D}).$$

As highlighted in the introduction of the paper, the proof of Theorem 2 will be via the surrogate quantity, $\mathrm{MenuGap}^2(X, Q)$. We will first show that for any distribution over two items $\mathcal{D}$, there exists two sequences of points (X, Q) such that $\mathrm{MenuGap}^2(X, Q)$ upper bounds the ratio between the revenue-optimal buy-two mechanism for $\mathcal{D}$ and the revenue from bundling, up to a constant factor (Lemma 9). Next we will show that this quantity itself is upper bounded for *all* pairs of sequences (X, Q) by another constant (Lemma 10). The proof of Theorem 2 will then follow directly from Lemmas 9, 10.

Lemma 9. *For any distribution $\mathcal{D}$ over 2 items, there exists a sequence of points $X = \{\vec{x}_i\}_{i\geq 1}$, $Q = \{\vec{q}_i\}_{i\geq 0}$ (starting with $\vec{q}_0 = (0, 0)$) such that*

$$\mathrm{MenuGap}^2(X, Q) \geq \frac{\mathrm{Buy2Rev}(\mathcal{D})}{20 \cdot \mathrm{BRev}(\mathcal{D})}.$$

The second piece of the proof is to upper bound this quantity itself.

Lemma 10. *For all sequences X, Q as defined in Definition 6 it holds that* $\mathrm{MenuGap}^2(X, Q) \leq 2$.

Proof of Theorem 2. Follows directly from Lemmas 9, 10. □

We devote the remainder of this section to the proofs of Lemmas 9, 10.

3.1 Proof of Lemma 9

The proof of Lemma 9 is split into two parts. In the first part, we will take the revenue-optimal buy-two menu for $\mathcal{D}$ and massage it down to a sub-menu of interest whose buy-two revenue remains close to the optimal one. The sub-menu itself may not be buy-two incentive compatible. The key lies in *only* removing menu entries, not altering existing ones. This does not change the incentive structure for the buyers whose options have remain intact. In the second part, we will show how to use that appropriate sub-menu in order to construct the desired sequence of points. The proof of Lemma 9 follows the blueprint of a similar lemma for the buy-one case from [25].

3.1.1 Finding a sub-menu of Interest

Let $\mathcal{M}^* = \{(p_i, \vec{q}_i)\}_{i=1}$ be the revenue-optimal buy-two incentive-compatible menu, where $(p_i, \vec{q}_i)$ denotes the price and expected allocation of the i-th entry of the menu.

Claim 11. *Let $\mathcal{M}$ be a buy-two incentive compatible mechanism, $\mathcal{M}_c \subseteq \mathcal{M}$ be the sub-menu of $\mathcal{M}$ that only offers options of price at least c. Then* $\text{Buy2Rev}(\mathcal{D}, \mathcal{M}_c) \geq \text{Buy2Rev}(\mathcal{D}, \mathcal{M}) - c$.

Proof. If a buyer with valuation $\vec{v}$ 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 $cf(\vec{v})$. Let S_c be the set of valuation vectors that preferred a menu entry in $\mathcal{M}$ priced at $p < c$. Then the total lost in revenue is at most $c \sum_{\vec{v} \in S_c} f(\vec{v}) \leq c$. □

Claim 12. *Let $\mathcal{M}$ be a mechanism whose prices are all at least c, 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 menu entries whose price $p_i \in \cup_{i=0}^{\infty} [c \cdot 3^{2i}, c \cdot 3^{2i+1})$.*
- *$\mathcal{M}_2$ has all menu entries whose price $p_i \in \cup_{i=0}^{\infty} [c \cdot 3^{2i+1}, c \cdot 3^{2i+2})$.*

Then $\max_{i=1,2} \text{Buy2Rev}(\mathcal{D}, \mathcal{M}_i) \geq \text{Buy2Rev}(\mathcal{D}, \mathcal{M})/2$.

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

$$\text{Buy2Rev}(\mathcal{D}, \mathcal{M}) \leq \text{Buy2Rev}(\mathcal{D}, \mathcal{M}_1) + \text{Buy2Rev}(\mathcal{D}, \mathcal{M}_2).$$

This is because any buyer with valuation $\vec{v}$ 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. □

Lemma 13. *There exists a menu $\mathcal{M}$ such that*

- *All prices are at least c.*
- *All prices belong to the union of intervals $\bigcup_{i=0}^{\infty} [c \cdot 3^{2i+a}, c \cdot 3^{2i+a+1})$ for an $a \in \{0, 1\}$.*
- $\text{Buy2Rev}(\mathcal{D}, \mathcal{M}) \geq \frac{\text{Buy2Rev}(\mathcal{D}) - c}{2}$.

Proof. Take the revenue-optimal buy-two menu $\mathcal{M}^*$, apply Claim 11 to obtain a menu $\mathcal{M}'$ that satisfies the first bullet point. Take the menu $\mathcal{M}'$ and apply Claim 12 to obtain a menu $\mathcal{M}$ that immediately satisfies the first and second bullet points. Finally, due to the revenue guarantees of Claims 11, 12, we have that $\text{Buy2Rev}(\mathcal{D}, \mathcal{M}) \geq \frac{\text{Buy2Rev}(\mathcal{D}, \mathcal{M}')}{2} \geq \frac{\text{Buy2Rev}(\mathcal{D}) - c}{2}$. □

We will eventually choose $c = O(\text{Buy2Rev}(\mathcal{D}))$. Thus at this point we have shown that there will always exist a somewhat structured sub-menu whose buy-two revenue is approximately optimal. The menu itself may not be buy-two incentive compatible.

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 $\text{MenuGap}^2(X, Q) \geq O\left(\text{Buy2Rev}(\mathcal{D})/\text{BRev}(\mathcal{D})\right)$.

Consider the menu $\mathcal{M}$ from Lemma 13. Let $\mathcal{B}_j \subseteq \mathcal{M}$ be the sub-menu that has all menu entries priced in $[3^{2j+a}, 3^{2j+a+1})$ for the same $a \in \{0, 1\}$ from Lemma 13. Let $\vec{x}_j$ be a valuation on $\mathcal{B}_j$ such that $|\vec{x}_j|_1 \leq (1+\delta)|\vec{x}|_1$ for all $\vec{x} \in \mathcal{B}_j$. We call such a valuation the *representative of bin* $\mathcal{B}_j$. Let X be the sequence defined by representatives and Q be the sequence defined by the respective allocations $\vec{q}_j \in \mathcal{M}$ of each representative $\vec{x}_j$, with $\vec{q}_0 = \vec{0}$.

We now prove a series of claims that will allow us to prove the main result of this subsection. The first claim simply upper bounds the probability of sampling a vector from each bin as a function of the norm of the representative and $\text{BRev}(\mathcal{D})$, the benchmark of interest.

Claim 14. $\Pr(\vec{x} \in \mathcal{B}_j) \leq \frac{\text{BRev}(\mathcal{D})(1+\delta)}{|\vec{x}_j|_1}$.

Proof. Consider the mechanism that sells the grand bundle at price $|\vec{x}_j|_1/(1+\delta)$. 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 $\Pr(\vec{x} \in \mathcal{B}_j)|\vec{x}_j|_1/(1+\delta)$. But this is a grand bundling menu and is thus its revenue is at most $\text{BRev}(\mathcal{D})$. □

The next claim lower bounds the gap contribution of the j-th representative as a function of parameters from Claim 14 as well as the price p_j of the menu entry that allocates $\vec{q}_j$.

Claim 15. $\text{gap}_j^2(X, Q) \geq \frac{\Pr(\vec{x} \in \mathcal{B}_j)p_j}{3 \cdot \text{BRev}(\mathcal{D})(1+\delta)}$.

Proof. Because the initial mechanism $\mathcal{M}^*$ was buy-two incentive-compatible, we know that for any pair of options $\vec{q}_{j'}, \vec{q}_{j''}$,

$$\vec{x}_j \cdot \vec{q}_j - p_j \geq \vec{x}_j \cdot \vec{\text{Lot}}(\vec{q}_{j'}, \vec{q}_{j''}) - p_{j'} - p_{j''}.$$

We can rewrite this as

$$\text{gap}_j^2(X, Q) \geq \frac{p_j - p_{j'} - p_{j''}}{|\vec{x}_j|_1} \geq \frac{p_j}{(3 \cdot |\vec{x}_j|_1)} \geq \frac{\Pr(\vec{x} \in \mathcal{B}_j)p_j}{3 \cdot \text{BRev}(\mathcal{D})(1+\delta)}.$$

Recall by our choice of points and the fact that $j', j'' < j$, $p_j \geq 3 \cdot p_{j'}, 3 \cdot p_{j''}$. Thus, the second inequality follows. The third inequality follows from Claim 14 since we have that $\frac{1}{|\vec{x}_j|_1} \geq \frac{\Pr(\vec{x} \in \mathcal{B}_j)p_j}{\text{BRev}(\mathcal{D})(1+\delta)}$.

□

We are now ready to prove Lemma 9.

Proof of Lemma 9. Let us first observe the following fact

$$\text{Buy2Rev}(\mathcal{D}, \mathcal{M}) = \sum_j \sum_{\vec{x} \in \mathcal{B}_j} p(x) f(x) \leq \sum_j \Pr(\vec{x} \in \mathcal{B}_j) 3p_j. \tag{3}$$

Recall that the price any valuations $\vec{x} \in \mathcal{B}_j$, its price $p(\vec{x})$ is no greater than $3p_j$. Therefore, the inequality follows. Moreover, from Claim 15 we get that

$$\text{MenuGap}^2(X, Q) = \sum_j \text{gap}_j^2(X, Q) \geq \sum_j \frac{\Pr(\vec{x} \in \mathcal{B}_j) p_j}{3 \cdot \text{BRev}(\mathcal{D})(1 + \delta)}. \tag{4}$$

Applying Eqn. 3 together with Lemma 13 we get that

$$\sum_j \Pr(\vec{x} \in \mathcal{B}_j) p_j \geq \text{Buy2Rev}(\mathcal{D}, \mathcal{M})/3 \geq (\text{Buy2Rev}(\mathcal{D},) - c)/6. \tag{5}$$

Putting Eqns. 4, 5 we get that

$$\text{MenuGap}^2(X, Q) \geq \frac{(\text{Buy2Rev}(\mathcal{D}) - c)}{18 \cdot \text{BRev}(\mathcal{D})(1 + \delta)}. \tag{6}$$

Let $c = \text{Buy2Rev}(\mathcal{D})/100, \delta = 1/100$ in Eqn. 6. Therefore,

$$\text{MenuGap}^2(X, Q) \geq \frac{99 \cdot \text{Buy2Rev}(\mathcal{D})}{101 \cdot 18 \cdot \text{BRev}(\mathcal{D})} \geq \frac{\text{Buy2Rev}(\mathcal{D})}{20 \cdot \text{BRev}(\mathcal{D})}.$$

□

3.2 Proof of Lemma 10

Given a sequence of points Q, let Q_i be the sequence truncated at the i-th point, that is to say $Q_i = \{\vec{q}_0, \vec{q}_1, ..., \vec{q}_i\}$. Let $\vec{m}_i = (\max_{\vec{q}_i \in Q_i}\{\vec{q}_{i,1}\}, \max_{\vec{q}_i \in Q_i}\{\vec{q}_{i,2}\})$ be the 2-dimensional vector whose entries are the largest coordinates among the points in Q_i.

Claim 16. *For any i,* $\text{gap}_i^2(X, Q)/||\vec{x}_i||_1 \leq \max\{\vec{q}_{i,1} - \vec{m}_{i-1,1}, 0\} + \max\{\vec{q}_{i,2} - \vec{m}_{i-1,2}, 0\}$.

Proof. Let $\vec{x}_i, \vec{q}_i$ be given. Since $\text{gap}_i^2(X, Q)$ is defined to be the minimum over all pairs of previously placed points, we can just upper bound it by witnessing its value with two earlier points. Let i_1^*, i_2^* be the indices such that:

- $i_1^*, i_2^* < i$,
- $\vec{q}_{i_1^*,1} = \vec{m}_{i-1,1}$,
- $\vec{q}_{i_2^*,2} = \vec{m}_{i-1,2}$.

That is to say, i_1^*, i_2^* are the indices of the points that witness that $\vec{m}_{i-1}$ is indeed the coordinate-wise maximum of all points in Q_i. Then

$$\frac{\mathrm{gap}_i^2(X,Q)}{||\vec{x}_i||_1} \leq \frac{\vec{x}_i}{||\vec{x}_i||_1} \cdot (\vec{q}_i - \vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*})). \tag{7}$$

Recall that

$$\begin{aligned}
\vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*})_1 &\geq \max\{\vec{q}_{i_1^*,1}, \vec{q}_{i_2^*,1}\} \geq \vec{m}_{i-1,1},\\
\vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*})_2 &\geq \max\{\vec{q}_{i_1^*,2}, \vec{q}_{i_2^*,2}\} \geq \vec{m}_{i-1,2}.
\end{aligned}$$

Therefore,

$$\vec{q}_{i,1} - \vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*}) \leq \vec{q}_{i,1} - \vec{m}_{i-1,1}, \tag{8}$$

$$\vec{q}_{i,2} - \vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*}) \leq \vec{q}_{i,2} - \vec{m}_{i-1,2}. \tag{9}$$

Therefore, for any choice of $\vec{x}_i$, it will be true that

$$\begin{aligned}
\frac{\mathrm{gap}_i^2(X,Q)}{||\vec{x}_i||_1} &\leq \frac{\vec{x}_i}{||\vec{x}_i||_1} \cdot (\vec{q}_i - \vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*}))\\
&= \frac{\vec{x}_{i,1}}{||\vec{x}_i||_1}(\vec{q}_{i,1} - \vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*})_1) + \frac{\vec{x}_{i,2}}{||\vec{x}_i||_1}(\vec{q}_{i,2} - \vec{\mathrm{Lot}}(\vec{q}_{i_1^*}, \vec{q}_{i_2^*})_2)\\
&\leq \frac{\vec{x}_{i,1}}{||\vec{x}_i||_1}(\vec{q}_{i,1} - \vec{m}_{i-1,1}) + \frac{\vec{x}_{i,2}}{||\vec{x}_i||_1}(\vec{q}_{i,2} - \vec{m}_{i-1,2})\\
&\leq \max\{0, \vec{q}_{i,1} - \vec{m}_{i-1,1}\} + \max\{0, \vec{q}_{i,2} - \vec{m}_{i-1,2}\}.
\end{aligned}$$

□

With Claim 16 at hand, the proof of Lemma 10 will follow easily.

Proof of Lemma 10. For any pair of sequences (X, Q),

$$\begin{aligned}
\mathrm{MenuGap}^2(X,Q) &= \sum_{i=1}^{N} \mathrm{gap}_i^2(X,Q)/||x_i||_1\\
&\leq \sum_{i=1}^{N}(\max\{\vec{q}_{i,1} - \vec{m}_{i-1,1}, 0\}) + \sum_{i=1}^{N}(\max\{\vec{q}_{i,2} - \vec{m}_{i-1,2}, 0\})\\
&\leq \sum_{i=1}^{N}(\vec{m}_{i,1} - \vec{m}_{i-1,1}) + \sum_{i=1}^{N}(\vec{m}_{i,2} - \vec{m}_{i-1,2})\\
&\leq 2.
\end{aligned}$$

The first inequality follows from Claim 16. For the second inequality, first note that by definition $\vec{q}_{i,1} \leq \vec{m}_{i,1}$, with equality only if $\vec{q}_{i,1} \geq \vec{q}_{i',1}$ for all $i' < i$. But note also that by definition $\vec{m}_{i,1} - \vec{m}_{i-1,1} \geq 0$. Therefore $\max\{\vec{m}_{i,1} - \vec{m}_{i-1,1}, 0\} \leq \vec{m}_{i,1} - \vec{m}_{i-1,1}$. The last inequality follows from observing that the final sum across each coordinate telescopes. Since $\vec{q}_{i,1} \leq 1$, the sum is at most 1 per coordinate.

□

Observation 17. *Setting both sequences (X, Q) equal to the standard basis of $\mathbb{R}^2$ shows that Lemma 10 is tight.*

4 Proof of Theorem 3

The proof of Theorem 3 follows an outline similar to that of Theorem 2. In this section we present only an outline of the proof, highlighting the relevant changes. All proofs that are sufficiently similar will be deferred to Appendix A or their corresponding proof in Section 3. We restate Theorem 3 for convenience.

Theorem 3. *For any distribution $\mathcal{D}$ over n items, it holds that*

$$O(n^3) \cdot \mathrm{BRev}(\mathcal{D}) \geq \mathrm{Buy}(n)\mathrm{Rev}(\mathcal{D}).$$

The proof of Theorem 3 will be via a surrogate quantity, $\mathrm{MenuGap}^k(X, Q)$, a natural generalization of $\mathrm{MenuGap}^2(X, Q)$ for buy-k mechanisms. We will first show that for any distribution $\mathcal{D}$ over n items, there exists two sequences of points such that $\mathrm{MenuGap}^k(X, Q)$ upper bounds the ratio between the revenue-optimal buy-k mechanism for $\mathcal{D}$ and the revenue from bundling, up to a $\text{POLY}(k)$ factor (Lemma 18). Next we will show that this quantity itself is upper bounded for *all* pairs of sequences (X, Q) for the case where $k = n$ (Lemma 19). The proof of Theorem 3 will follow directly from Lemmas 18, 19.

Lemma 18. *For any distribution $\mathcal{D}$ over n items, there exists a sequence of points $X = \{\vec{x}_i\}_{i=1}$, $Q = \{\vec{q}_i\}_{i=0}$ (starting with $\vec{q}_0 = (0, \ldots, 0)$) such that*

$$\mathrm{MenuGap}^k(X, Q) \geq \frac{\mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D})}{3(k+1)^2 \cdot \mathrm{BRev}(\mathcal{D})}.$$

Note that in Lemma 18 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 the analog of Lemma 10 for the case where $k \geq n$.

Lemma 19. *For all sequences X, Q as defined in Definition 6 it holds that $\mathrm{MenuGap}^n(X, Q) \leq n$.*

Proof of Theorem 3. Follows directly from Lemma 18 (setting $k = n$) and Lemma 19. □

4.1 Proof of Lemma 18

The proof of Lemma 18 is split into two parts. In the first part, we will take the revenue-optimal buy-k menu for $\mathcal{D}$ and massage it down to a sub-menu of interest whose revenue remains close to the optimal one. The sub-menu itself may not be buy-k incentive-compatible. However, similar to Lemma 9, the key to approximately preserving the revenue will be in just removing entries from the revenue-optimal 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 18 follows the blueprint of Lemma 9.

4.1.1 Finding a Sub-menu of Interest

Let $\mathcal{M}^* = \{(p_i, \vec{q}_i)\}_{i=1}$ be the revenue-optimal buy-k, where $(p_i, \vec{q}_i)$ denotes the price and expected allocation of the i-th option of the menu.

Claim 20. *Let $\mathcal{M}$ be a buy-two incentive-compatible mechanism, $\mathcal{M}_c \subseteq \mathcal{M}$ 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. Identical to the proof of Claim 11. □

Claim 21. *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 $p_i \in \cup_{i=0}^{\infty}[c \cdot (k+1)^{2i}, c \cdot (k+1)^{2i+1})$.*
- *$\mathcal{M}_2$ has all options whose price $p_i \in \cup_{i=0}^{\infty}[c \cdot (k+1)^{2i+1}, c \cdot (k+1)^{2i+2})$.*

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. Identical to the proof of Claim 12. □

Lemma 22. *There exists a menu $\mathcal{M}$ such that*

- *All prices are at least c.*
- *All prices belong to the set of intervals $\bigcup_{i=0}^{\infty}[c\cdot(k+1)^{2i+a}, c\cdot(k+1)^{2i+a+1})$ 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 revenue-optimal buy-k menu $\mathcal{M}^*$, apply Claim 20 to obtain a menu $\mathcal{M}'$ that satisfies the first bullet point. Take the menu $\mathcal{M}'$ and apply Claim 21 to obtain a menu $\mathcal{M}$ that immediately satisfies the first and second bullet points. Finally, due to the revenue guarantees of Claims 20, 21, we have that $\text{Buy}(k)\text{Rev}(\mathcal{D}, \mathcal{M}) \geq \frac{\text{Buy}(k)\text{Rev}(\mathcal{D},\mathcal{M}')}{2} \geq \frac{\text{Buy}(k)\text{Rev}(\mathcal{D})-c}{2}$. □

4.1.2 Construction 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 $\text{MenuGap}^k(X, Q) \geq O(\text{Buy}(k)\text{Rev}(\mathcal{D})/\text{BRev}(\mathcal{D}))$. Consider the menu $\mathcal{M}$ from Lemma 22. Let $\mathcal{B}_j \subseteq \mathcal{M}$ be the sub-menu that has all menu entries priced in $[c \cdot (k+1)^{2j+a}, c \cdot (k+1)^{2j+a+1})$ for the same $a \in \{0, 1\}$ from Lemma 22. Let $\vec{x}_j$ be the valuation on $\mathcal{B}_j$ such that $|\vec{x}_j|_1 \leq (1+\delta)|\vec{x}|_1 \forall \vec{x} \in \mathcal{B}_j$. We call vector $\vec{x}_j$ the *representative of bin $\mathcal{B}_j$*.

Claim 23. $\Pr(\vec{x} \in \mathcal{B}_j) \leq \frac{\text{BRev}(\mathcal{D})(1+\delta)}{|\vec{x}_j|_1}$.

Proof. Identical to Claim 14. □

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

Claim 24. $\text{gap}_j^k(X, Q) \geq \frac{p_j}{(k+1)\cdot|\vec{x}_j|_1} \geq \frac{\Pr(\vec{x}\in\mathcal{B}_j)p_j}{(k+1)\cdot\text{BRev}(\mathcal{D})(1+\delta)}$.

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

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

We can rewrite this as

$$\text{gap}_j^k(X, Q) \geq \frac{p_j - \sum_{i=1}^k p_{j_i}}{|\vec{x}_j|_1}.$$

Recall by our choice of points and the fact that $j_i < j$, $p_j \geq (k+1)p_{j_i}$. Therefore, the right hand is at least $\frac{p_j}{(k+1)|\vec{x}_j|_1}$.

The second inequality comes from Claim 23.

□

We are now ready to prove Lemma 18.

Proof of Lemma 18. Let us first observe that

$$\text{Buy}(k)\text{Rev}(\mathcal{D}, \mathcal{M}) = \sum_j \sum_{\vec{x}\in\mathcal{B}_j} p(\vec{x})f(\vec{x}) \leq \sum_j \Pr(\vec{x} \in \mathcal{B}_j)(k+1)p_j. \tag{10}$$

Recall that the price any valuations $\vec{x} \in \mathcal{B}_j$, its price $p(\vec{x})$ is no greater than $(k+1)p_j$. Therefore, the inequality follows. Moreover, from Claim 24 we get that

$$\text{MenuGap}^k(X, Q) = \sum_j \text{gap}_j^k(X, Q) \geq \sum_j \frac{\Pr(\vec{x} \in \mathcal{B}_j)p_j}{(k+1)\cdot\text{BRev}(\mathcal{D})(1+\delta)}. \tag{11}$$

Applying Eq. 10 together with Lemma 22 we get that

$$\sum_j \Pr(\vec{x} \in \mathcal{B}_j)p_j \geq \text{Buy}(k)\text{Rev}(\mathcal{D}, \mathcal{M})/(k+1) \geq \frac{(\text{Buy}(k)\text{Rev}(\mathcal{D}) - c)}{2(k+1)}. \tag{12}$$

Putting Eqs. 11, 12 we get that

$$\text{MenuGap}^k(X, Q) \geq \frac{(\text{Buy}(k)\text{Rev}(\mathcal{D}) - c)}{2(k+1)^2 \cdot \text{BRev}(\mathcal{D})(1+\delta)}. \tag{13}$$

Let $c = \text{Buy2Rev}(\mathcal{D})/100, \delta = 1/100$ in Equation 13. Therefore,

$$\frac{99 \cdot \text{BuyManyRev}(\mathcal{D})}{101 \cdot 2(k+1)^2 \cdot \text{BRev}(\mathcal{D})} \geq \frac{\text{BuyManyRev}(\mathcal{D})}{3 \cdot (k+1)^2 \cdot \text{BRev}(\mathcal{D})}.$$

□

4.2 Proof of Lemma 19

The proof of Lemma 19 will be similar to that of Lemma 10. Given a sequence of points Q, let Q_i be the sequence truncated at the i-th point, that is to say $Q_i = \{\vec{q}_0, \vec{q}_1, \ldots, \vec{q}_i\}$. Let $\vec{m}_i = (\max_{\vec{q}_i \in Q_i}\{\vec{q}_{i,1}\}, \ldots, \max_{\vec{q}_i \in Q_i}\{\vec{q}_{i,n}\})$ be the n-dimensional vector whose entries are the largest coordinates among the points in Q_i.

Claim 25. *For any i, $\mathrm{gap}_i^n(X, Q)/||\vec{x}_i||_1 \leq \sum_{d=1}^{n} \max\{\vec{q}_{i,d} - \vec{m}_{i-1,d}, 0\}$.*

Proof. Similar to that of Claim 16. Deferred to Appendix A. □

Proof of Lemma 19. For any pair of sequences (X, Q),

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

The first inequality follows from Claim 25. 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 $\vec{m}_{i,1} - \vec{m}_{i-1,1} \geq 0$. Therefore $\max\{\vec{m}_{i,1} - \vec{m}_{i-1,1}, 0\} \leq \vec{m}_{i,1} - \vec{m}_{i-1,1}$. 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.

□

Observation 26. *Setting both sequences (X, Q) equal to the standard basis of $\mathbb{R}^n$ shows that Lemma 19 is tight.*

5 Proof of Theorem 4

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 $\mathrm{POLY}(n)$) between $\mathrm{BRev}(\mathcal{D})$ and $\mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D})$, the revenue attained by the optimal buy-k mechanism.

Theorem 4. *If $k \leq n^{1/2-\varepsilon}$ for some $\varepsilon > 0$, then there exists a distribution $\mathcal{D}$ over n items such that*

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

The proof of Theorem 4 will be broken down in three steps. Firstly, we will describe the pair of sequences (X^L, Q^L) that we use (Subsection 5.1.1). The construction will make use of a combinatorial Lemma about cover-free sets from [20]. 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 32). 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 33). The proof of Lemma 33 will use ideas from [18].

Before we delve into the proof of Theorem 4, we analyze its implications for mechanisms with polynomial menu size. We invoke the following Corollary from [18].

Corollary 27 (Restated from [18])**.** *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 4 imply the following Corollary.

Corollary 28. *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*

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

In other words, Corollary 28 rules out all polynomial-sized mechanisms $\mathcal{M}$ as candidates for good approximations to $\text{Buy}(k)\text{Rev}(\mathcal{D})$ for the case $k \leq n^{1/2-\varepsilon}$.

5.1 Proof of Theorem 4

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

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 [15], attributed there to [20].

Theorem 30 ([20])**.** *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 $|\mathcal{F}_k| \geq 2^{\Omega\left(\frac{n}{k^2}\right)}$.

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 31. *Let $\mathcal{F}^L$ be a k-cover-free family of sets of maximal size, i.e., such that $|\mathcal{F}^L| = T(n,k) = \exp\left(\Omega(n/k^2)\right)$. 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., [18], [25]) 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.

5.1.2 Part 2: The Instance has Large Menu Gap

In the next step of the proof of Theorem 4 we will show that the constructed instance has large menu gap.

Claim 32. *For the instance described in Definition 31, it holds that*

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

Proof.

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

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 $A_i \setminus \left(\cup_{\ell=1}^k A_{j_\ell}\right) \neq \emptyset$ 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, $\text{MenuGap}^k(X^L, Q^L) \geq \frac{|\mathcal{F}^L|}{n}$. $\square$

5.1.3 Part 3: from Sequences to Distributions

We now present the final piece for the proof of Theorem 4. Lemma 33 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 [18]. 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 33. *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 $\vec{v}$ and a density function $f(\vec{v})$. Let $C_i = (n+1)^{2i}$. Then we define distribution $\mathcal{D}$ by setting $\vec{v}_i = \vec{x}_i \cdot C_i$, $f(\vec{v}_i) = \frac{1}{C^i}$. It is clear that this defines a valid distribution, i.e., $f(\vec{v}_i) \geq 0$ and $\sum_i f(\vec{v}_i) \leq 1$. 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 34. *Any valuation $\vec{v}_i$ 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 $\vec{v}_i$ 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 $\vec{v}_i$ 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 $\vec{v}_i \cdot \vec{\text{Lot}}(\vec{q}_{i_1}, \vec{q}_{i_2}, \ldots, \vec{q}_{i_k})$. The utility from purchasing $(p_i, \vec{q}_i)$ is $\vec{v}_i \cdot \vec{q}_i - p_i$. By choice of p_i, $\vec{v}_i$ we get that this is

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

where the inequality follows from recalling the definition of $\text{gap}_k^i(X, Q)$ (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 $\vec{v}_i$ will purchase exactly one copy of the menu entry $(p_i, \vec{q}_i)$. The revenue of mechanism $\mathcal{M}$ is $\sum_i f(\vec{v}_i)p_i = \sum_i \text{gap}_i^k(X, Q) = \text{MenuGap}^k(X, Q)$. 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 $\vec{v}_i$ 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 $b_i \cdot \Pr_{\vec{v}_j \sim \mathcal{D}}(C_j|\vec{x}_j| \geq b_i) = b_i \cdot \Pr_{\vec{v}_j \sim \mathcal{D}}(C_j \geq C_i) = b_i \cdot \sum_{j \geq i} f(\vec{v}_j) = b_i \cdot \frac{2n}{C_i(2n-1)} \leq 2b_i \cdot \frac{1}{C_i} \leq 2n$. □

5.1.4 Part 4: Putting it all together

We are now ready to present the Proof of Theorem 4.

Proof of Theorem 4. Consider the instance (X^L, Q^L) from Definition 31. By Claim 32, the instance satisfies $\text{MenuGap}^k(X^L, Q^L) \geq |\mathcal{F}^L|/n$ and has only integral vectors. By Lemma 33, 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 $\frac{\exp(\Omega(n^\varepsilon))}{2n^2}$.

□

6 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 gap between buy-many and 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 $O(n^3)$ fraction of the optimal buy-k revenue. Our proof uses a recent framework proposed in [18, 25] 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. Firstly, it would be interesting to understand whether or not Lemma 18 is tight. Any improvements to the approximation ratio would directly translate to improvements on Theorem 3. As observed earlier in the paper, Lemma 19 is tight. In addition, our techniques seem promising for the task of extending our results beyond additive valuations.

Another important question is 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 $\mathrm{Buy}(k)\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D})$ 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 [13] 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, computationally very little is still understood about buy-many or 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 Addi-

tive 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 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

[6] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*.

[7] 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)*.

[8] 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

[9] 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

[10] Shuchi Chawla, Rojin Rezvan, Yifeng Teng, and Christos Tzamos. 2021. Pricing Ordered Items. *CoRR* abs/2106.04704 (2021). arXiv:2106.04704 https://arxiv.org/abs/2106.04704

[11] 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* (Phoenix, AZ, USA) *(EC '19)*. Association for Computing Machinery, New York, NY, USA, 237–238. https://doi.org/10.1145/3328526.3329583

[12] 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

[13] 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* (Virtual Event, Hungary) *(EC '20)*. Association for Computing Machinery, New York, NY, USA, 475–476. https://doi.org/10.1145/3391403.3399453

[14] Constantinos Daskalakis. 2015. Multi-Item Auctions Defying Intuition? *SIGecom Exch.* 14, 1 (nov 2015), 41–75. https://doi.org/10.1145/2845926.2845928

[15] 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

[16] Sergiu Hart and Noam Nisan. 2012. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*.

[17] Sergiu Hart and Noam Nisan. 2013. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*.

[18] 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

[19] Sergiu Hart and Philip J. Reny. 2015. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics* 10, 3 (2015), 893–922.

[20] 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

[21] Pravesh Kothari, Sahil Singla, Divyarthi Mohan, Ariel Schvartzman, 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

[22] 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.

[23] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.

[24] Alexandros Psomas, Ariel Schvartzman, 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*

(Phoenix, AZ, USA) *(EC '19)*. Association for Computing Machinery, New York, NY, USA, 417–418. https://doi.org/10.1145/3328526.3329563

[25] Christos-Alexandros Psomas, Ariel Schvartzman, and S. Matthew Weinberg. 2022. On Infinite Separations Between Simple and Optimal Mechanisms: A Converse of a Theorem of Hart and Nisan. *Manuscript* (2022).

[26] 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.*

[27] 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 Proofs Omitted from Section 4

Proof of Claim 25. 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 it's value with two earlier points. Let $i_1^*, i_2^*, \ldots, i_n^*$ be the indices such that:

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

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 placed before i. Then

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

Recall that

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

Therefore,

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

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

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

This proves the claim (Naturally, $\vec{x}_{i,1} \le ||\vec{x}_i||_1, \vec{x}_{i,2} \le ||\vec{x}_i||_1$).

□

B 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 [11].

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 $\vec{\mathrm{Lot}}(\cdot)$.

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 [11], 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 [11] for more details on this definition).

As was observed in [11], 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 2 and 3 to adaptive buy-k incentive-compatible mechanisms.

Finally recall that the construction of Theorem 4 presented in Section 5 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 4 also extend to the case of adaptive buyers.