

A Multi-Dimensional Online Contention Resolution Scheme for Revenue Maximization*

Shuchi Chawla
shuchi@cs.utexas.edu

Dimitris Christou
christou@cs.utexas.edu

Trung Dang
dddtrung@cs.utexas.edu

Zhiyi Huang
zhiyih@cs.utexas.edu

Gregory Kehne
gkehne@utexas.edu

Rojin Rezvan
rojin@cs.utexas.edu[†]

Abstract

We study multi-buyer multi-item sequential item pricing mechanisms for revenue maximization with the goal of approximating a natural fractional relaxation – the ex ante optimal revenue. We assume that buyers’ values are subadditive but make no assumptions on the value distributions. While the optimal revenue, and therefore also the ex ante benchmark, is inapproximable by any simple mechanism in this context, previous work has shown that a weaker benchmark that optimizes over so-called “buy-many” mechanisms can be approximable. Approximations are known, in particular, for settings with either a single buyer or many unit-demand buyers. We extend these results to the much broader setting of many subadditive buyers. We show that the ex ante buy-many revenue can be approximated via sequential item pricings to within an $O(\log^2 m)$ factor, where m is the number of items. We also show that a logarithmic dependence on m is necessary.

Our approximation is achieved through the construction of a new multi-dimensional Online Contention Resolution Scheme (OCRS), that provides an online rounding of the optimal ex ante solution. Chawla et al. [2023] previously constructed an OCRS for revenue for unit-demand buyers, but their construction relied heavily on the “almost single dimensional” nature of unit-demand values. Prior to that work, OCRSes have only been studied in the context of social welfare maximization for single-parameter buyers. For the welfare objective, constant-factor approximations have been demonstrated for a wide range of combinatorial constraints on item allocations and classes of buyer valuation functions. Our work opens up the possibility of a similar success story for revenue maximization.

*This work was funded in part by NSF award CCF-2217069.

[†]All co-authors are affiliated with the University of Texas at Austin.

1 Introduction

In the last decade, strong connections have emerged between mechanism design, online selection, and optimal stopping problems such as prophet inequalities, leading to a number of beautiful results and techniques in both areas. Among the surprising insights to emerge from this body of work is the small gap between online and offline optimization, as well as between simple and optimal mechanisms, for the objective of social welfare maximization. Consider, in particular, a seller wishing to allocate m items across n buyers with combinatorial valuations over the items, denoted by v_i for $i \in [n]$. The seller’s goal is to find a partition $\{S_1, S_2, \dots, S_n\}$ of the set $[m]$ of items and allocate set S_i to each buyer i so as to maximize the social welfare, defined as $\sum_{i \in [n]} v_i(S_i)$. The seller can do so using the VCG mechanism. However, from a single buyer’s perspective, his allocation and payment to the seller as functions of his and others’ values can be complicated and inscrutable. It turns out that the seller can use a much simpler mechanism at a small constant factor loss, when the buyers’ values are structured and drawn from known distributions. For example, if buyers have fractional subadditive (i.e. XOS) values, the seller can obtain a 2-approximation via a *sequential item pricing*: she interacts with each buyer one at a time, offering any remaining items at pre-computed fixed prices, and allowing the buyer to select his favorite set of items to buy.

But what if the seller wants to maximize her revenue? Can a simple online mechanism like sequential item pricing still approximate the optimal offline mechanism, and if so, under what assumptions? This work addresses and answers these questions.

In this paper, we make an explicit connection between revenue maximization for many combinatorial buyers and online contention resolution schemes (henceforth, OCRSes), a key technique in designing prophet inequalities. **We develop a framework for designing OCRSes for revenue maximization over a given class of mechanisms, and instantiate it with an OCRS over sequential item pricings.** As a consequence of this construction, we obtain novel simplicity versus optimality results for revenue maximization. Our results hold for arbitrary distributions over subadditive valuation functions. We discuss these implications for mechanism design in detail below.

Contention Resolution Schemes and their connections to sequential mechanisms

An OCRS is an online rounding of a fractional relaxation of the offline objective for packing-type problems. OCRSes were originally designed for maximizing set functions [Vondrák et al., 2011, Feldman et al., 2016], the natural algorithmic counterparts of “single parameter” prophet inequalities. Let us consider the example of a seller wishing to sell a single item. Each buyer’s type is described by a single value $v_i \sim \mathcal{D}_i$ for the item; the objective function is the expectation over the joint value distribution of the value of the buyer who receives the item. The offline optimum, a.k.a. the prophet, receives a reward of $\mathbb{E}_v[\max_i v_i]$. Let us now consider a fractional relaxation where $x_i \in [0, 1]$ denotes the probability that buyer i receives the item. Considering buyer i ’s contribution alone to the objective, this contribution is maximized if the buyer receives the item when his value is in the highest x_i quantiles of his distribution. Let $\text{SW}_{x_i}(\mathcal{D}_i)$ denote this expected contribution. The fractional relaxation then maximizes $\max_{x \in [0, 1]^n: \sum_i x_i \leq 1} \sum_i \text{SW}_{x_i}(\mathcal{D}_i)$. This fractional relaxation has a natural economic interpretation: we relax the *ex post* supply constraint (namely, that we have one item to sell) to an *ex ante* supply constraint (namely, that the expected number of items sold is one). Accordingly, this relaxation is called the *ex ante relaxation*.

An OCRS takes as input an *ex ante* feasible solution x and “rounds” it in an online fashion into an *ex post* feasible solution as follows: we first sample each $i \in [n]$ with probability x_i , forming a (possibly infeasible) set $R(x) \subseteq [n]$; we then pick a feasible subset of $R(x)$ in an online fashion. The goal is to show that (1) each element has a low probability of being “blocked” by previously selected elements, and (2) conditioned on being sampled and unblocked, each element is picked with sufficiently high probability by the online algorithm. Owing to the concavity of social welfare as a function of the allocation probability¹, we then obtain an approximation to the *ex ante* relaxation. Since the *ex ante* relaxation bounds the prophet’s reward from above, this immediately implies a prophet inequality. Furthermore, the online rounding can be

¹i.e., $\text{SW}_x(\mathcal{D}_i) \geq \frac{x}{y} \text{SW}_y(\mathcal{D}_i)$ for all $0 < x < y$.

interpreted as a sequential posted pricing mechanism as it effectively sets a value threshold for each buyer, above which the buyer is served.

Lee and Singla [2018] in fact showed a tight connection between prophet inequalities, OCRSes, and the ex ante relaxation: For any downwards closed feasibility constraint \mathcal{P} in the single parameter setting,² the optimal competitive ratio achievable against the prophet’s reward, which is also the worst-case gap between the optimal social welfare and that achievable by a sequential posted pricing, is *exactly* equal to the optimal competitive ratio achievable against the ex ante relaxation via an OCRS.

A multi-dimensional OCRS for revenue. We now turn to the revenue objective in the context of allocating many items. In a multi-dimensional OCRS, as in the single-dimensional case, our goal is to “round” an ex ante feasible solution x in an online/sequential fashion. Consider a setting with n buyers, where \mathcal{D}_i denotes buyer i ’s value distribution, and $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ denotes the joint value distribution across all buyers. Let x_i denote a vector of probabilities where the j th component x_{ij} is the probability of allocating item j to buyer i . We will say that a single-buyer mechanism for buyer i satisfies the ex ante constraint x_i if this mechanism allocates each item j with probability at most x_{ij} to the buyer, where the probability is taken over the randomness in the mechanism as well as the randomness in the buyer’s values. Then, we can write $\text{REV}_{x_i}(\mathcal{D}_i)$ as the maximum revenue that can be obtained from buyer i from any mechanism that satisfies the constraint x_i . The *ex ante relaxation for revenue* is then given by $\text{EAREV}(\mathcal{D}) := \max_x \sum_i \text{REV}_{x_i}(\mathcal{D}_i)$, where the maximum is taken over all vectors $\{x_i\}_{i \in [n]}$ satisfying $\sum_i x_i \preceq (1, 1, \dots, 1)$.³ Likewise, we can define $\text{REV}_{x_i}(\mathcal{D}_i, \mathcal{C})$ and $\text{EAREV}(\mathcal{D}, \mathcal{C})$ by optimizing over a specific class \mathcal{C} of single-buyer mechanisms.

The OCRS generates a sequential mechanism in which the seller interacts with the buyers one at a time. At each iteration i , the seller offers a (potentially different) single-buyer mechanism over the as-yet-unsold items, say S_i , to the current buyer i , and buyer i purchases some subset of these items. The goal is to design a single-buyer mechanism for each i that extracts a sufficiently large fraction of the target revenue $\text{REV}_{x_i}(\mathcal{D}_i)$ from S_i while continuing to satisfy the ex ante constraint x_i . Of course, in order for this to be possible, S_i should contain sufficiently many items. We formally define a revenue OCRS as follows.

Definition. An α -OCRS for revenue maximization under a class of mechanisms \mathcal{C} takes as input an ex ante allocation constraint $x \in [0, 1]^m$, a value distribution \mathcal{D} , and a random subset of items $S \subseteq [m]$, and returns a single-buyer mechanism $M(\mathcal{D}, S, x) \in \mathcal{C}$ such that:

- (a) M only allocates items in S to the buyer. Furthermore, for each $j \in S$, M allocates j to a buyer with value distribution \mathcal{D} with probability at most x_j .
- (b) M ’s expected revenue from a buyer with value distribution \mathcal{D} is at least a $1/\alpha$ fraction of $\beta \cdot \text{REV}_x(\mathcal{D}, \mathcal{C})$, where $\beta := \min_{j \in [m]} \Pr[j \in S]$.

We show that if a class \mathcal{C} of single-buyer mechanisms admits an α -OCRS, then there exists a sequential \mathcal{C} mechanism that achieves an $O(\alpha)$ -approximation to $\text{EAREV}(\mathcal{D}, \mathcal{C})$. In other words, the gap between online and offline \mathcal{C} mechanisms is bounded by $O(\alpha)$.

The existence of an OCRS depends on the properties of the class \mathcal{C} of mechanisms being considered, as well as on the valuation functions of the buyers. Our main theorem shows the existence of an $O(\log m)$ -OCRS for item pricing mechanisms and subadditive valuations.

Informal Theorem 1. *There exists an $O(\log m)$ -OCRS for revenue maximization under item pricing mechanisms when all buyers have (arbitrary distributions over) subadditive values over m items and the ex ante constraint x satisfies $x_j \geq 1/\text{poly}(m)$ for all $j \in [m]$.*

We remark that although multi-item, i.e. combinatorial, prophet inequalities have been studied extensively for the social welfare objective (see, e.g., Feldman et al. [2015]), they do not use OCRSes and are based on other techniques. We review this literature in Section 9. On the other hand, whereas prophet

²Here \mathcal{P} is the convex hull of the incidence vectors of all subsets of buyers that can be feasibly served.

³For vectors y and z , we say $y \preceq z$ if for all coordinates i we have $y_i \leq z_i$.

inequalities designed for the social welfare objective in the single-parameter setting extend immediately to the revenue objective via virtual values, no such general connections exist between revenue and social welfare in the multi-item setting (some exceptional special cases are discussed in Section 9). This necessitates designing an OCRS specifically tailored to the revenue objective. To our knowledge, the only work prior to ours containing a multi-dimensional OCRS for revenue is the recent work of Chawla et al. [2023] that proves a 2-OCRS for unit-demand valuations – a special case of our result above.

Implications for revenue maximization

Revenue maximization for buyers with multi-dimensional values is notoriously challenging. Even for settings with just one buyer and two items, the optimal revenue is inapproximable to within any factor by a simple mechanism (indeed, by any mechanism with a finite description complexity) [Briest et al., 2015, Hart and Nisan, 2019]. A recent line of work initiated by Chawla et al. [2022b] overcomes this challenge by considering revenue maximization over so-called *Buy-Many Mechanisms*, that essentially restrict the seller to offering subadditive pricings over allocations to the buyer. Chawla et al. showed that the optimal buy-many revenue for a single buyer is approximable to within an $O(\log m)$ factor by item pricings for any distribution over buyer values, where m is the number of items. This result is tight.

Multi-buyer settings present a further challenge in the manner that buyers impose externalities upon one another (i.e. the potential loss in revenue from allocating an item to one buyer instead of another). In single item settings, we can characterize this externality using virtual values. But in multi-item settings the unwieldy structure of multi-dimensional incentive constraints disallows such a characterization. The *ex ante relaxation* was first proposed by Alaei [2014] to address precisely this challenge: it allows breaking up and reducing the multi-agent problem into its single-agent counterparts and approximating each one separately.

In a recent work, Chawla et al. [2023] brought these two lines of work together and defined an ex ante relaxation for multi-buyer buy-many mechanisms, $\text{EAREV}(\mathcal{D}, \text{BUYMANY})$. They showed that when every buyer has a unit-demand valuation function, the buy-many ex ante relaxation can be approximated by a sequential item pricing mechanism to within an $O(\log m)$ factor. They left open the question of designing an approximation for the buy-many ex ante relaxation for other classes of valuations. Our work resolves this open question by bounding this gap for arbitrary distributions over *subadditive values* to within a $O(\log^2 m)$ factor as a direct consequence of our OCRS construction.

Informal Theorem 2. *For settings with subadditive buyers over m items, sequential item pricing obtains at least an $O(\log m)$ fraction of the ex ante item pricing revenue and at least an $O(\log^2 m)$ fraction of the optimal multi-buyer buy-many revenue.*

We summarize our results on the approximation of the ex-ante item pricing revenue by sequential item pricings for different classes of value distributions in Table 1. All of our upper bounds extend to approximation against the optimal buy-many revenue at a loss of an additional $O(\log m)$ factor via the results of [Chawla et al., 2023].

	Upper Bounds	Lower Bounds	Reference
Unit-Demand	2	2	[Chawla et al., 2023]
Gross Substitutes	2	2	Theorem 5 and [Chawla et al., 2023]
XOS	$O(\log m)$	$\Omega(\log^{1/2} m)$	Theorems 2 and 4
Subadditive	$O(\log m)$	$\Omega(\log^{1/2} m)$	Theorems 2 and 4
Monotone	$\min\{n, 4m^2\}$	$\Omega(\min\{n, \sqrt{m}\})$	Theorems 18 and 19

Table 1: Summary of known results and our contributions for approximating $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$ by $\text{REV}(\mathcal{D}, \text{SEQITEMPRICING})$ for different families of buyer valuations.

Designing the OCRS for item pricings

Let us now consider the single-agent problem at the heart of our multi-dimensional OCRS. We are given a random subset S of items that contains every item with (say) $1/2$ probability and an ex ante supply constraint x . We want to extract (say) half as much revenue from the buyer over this subset of items, as the optimal mechanism (call it M^*) obtains from the buyer over the entire set of items, with both our mechanism and the optimal one satisfying the given ex ante constraint x . Let T^* denote the set of items a buyer purchases under M^* . If we could manage to sell $T^* \cap S$ to the buyer at the same prices as M^* , then we would be done. But the trouble with having only a subset of items available for sale is that the buyer may switch from wanting to buy $T^* \cap S$ to a completely different set of items, say T . This leads to two challenges: first T may not be able to generate as much revenue as T^* ; second, this switch may cause us to violate the ex ante supply constraint by selling some items with a higher probability than intended (e.g. the items in $T \setminus T^*$). We show how to resolve both of these issues for item pricings.

Our first main technical contribution is to show that if the buyer has subadditive values and we are allowed to ignore the ex ante constraint x , for any given set S , we can generate revenue from the items in S that is comparable (within a logarithmic factor) to the revenue M^* obtains from the same set S of items. We use the subadditivity of values to argue that the buyer obtains sufficiently high utility under M^* from the items in S . Then, using an approach developed by Chawla et al. [2022b], we scale up the prices in M^* to extract a fraction of this utility as revenue. This *Revenue Recovery Scheme* (RRS) (formally defined in Section 2) satisfies the revenue requirement of an OCRS as well as guarantees that the *total* number of items sold is not much larger than the number sold by M^* . However, it does not necessarily satisfy the given *per-item* allocation constraints.

Our second technical contribution fixes the per-item allocation constraints while losing only a constant factor in revenue. We make use of the fact that a revenue recovery scheme as described above exists for every subset T of items and only allocates items in T . To obtain an OCRS, we carefully choose a distribution over all subsets T of S and apply the RRS to the chosen subset to obtain a random pricing. Suppose, for example, that the RRS applied to the entire set S oversells some item j . Then, instead of deterministically choosing the pricing corresponding to S , with some appropriate probability we drop the item j , apply the RRS to $S \setminus \{j\}$, and choose the corresponding pricing. Having to drop items in this manner hurts our revenue guarantee. But importantly, the RRS cannot oversell too many items as it satisfies the ex ante constraint in aggregate over all items. Therefore it becomes possible to choose a random set T such that the total revenue contribution over T is large enough, while at the same time, the per-item allocation probabilities match those of the given ex ante constraint. We call this procedure the *Convex Hull Sampler* as it produces a random pricing in the convex hull of those given by the RRS at different sets $T \subseteq S$.

Unfortunately, the multiplicative loss in revenue for the RRS we design depends on the ratio of the maximum to the minimum price charged in M^* , and this dependence is necessary. This dependence can arise due to a potential long tail of the buyer's value distribution leading to exponentially large prices coupled with exponentially small allocation probabilities. Our final technical contribution is to eliminate this dependence in our approximation of ex ante revenue. In particular, as long as each ex ante constraint x_j is at least $\Omega(1/\text{poly } m)$, it becomes possible to obtain an $O(\log m)$ -OCRS for item pricings. This in turn provides the $O(\log m)$ approximation to the ex ante revenue for item pricings stated in Informal Theorem 2.

Finally, we remark that our OCRS and sequential item pricing are fully constructive and can be found in time polynomial in n , m , and the support sizes of the distributions \mathcal{D}_i , assuming: (1) We have access to the ex ante optimal solution x and the corresponding optimal (random) pricing p ; (2) We have access to a demand oracle that given an item pricing and a value function in the support of $\cup_i \mathcal{D}_i$ returns the set of items bought by the buyer with value v .

Overview of the rest of the paper

In Section 2 we formally introduce the ex ante benchmark, the OCRS for revenue maximization, and an outline of our approach towards designing an OCRS for item pricings. We present the four main parts of our upper bound construction in Sections 3 to 6. All our constructions are presented as existential results;

we address the computational aspects in Section 7. In Section 8 we provide an $\Omega(\sqrt{\log m})$ lower bound on the existence of an OCSR for item pricing over the class of XOS valuation functions. Our lower bound in fact applies to the gap between the ex ante item pricing relaxation and the revenue of *any* ex post feasible item pricing mechanism. We defer a discussion of gross substitutes and general valuations to the appendix. Section 9 discusses related work. We conclude and outline directions for future work in Section 10.

2 Definitions and an Outline of Our Construction

We consider the standard mechanism design setting with m items and n buyers with the objective of maximizing the total revenue of the seller. Buyers have combinatorial valuations over the items, $v_i : 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$ for $i \in [n]$, drawn from known independent distributions \mathcal{D}_i . We write $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ as the joint value distribution. Section A of the appendix describes different classes of valuations and Bayesian incentive compatible (IC) mechanisms for readers unfamiliar with mechanism design. Henceforth, we assume knowledge of these concepts.

For a buyer with valuation v and a single-buyer mechanism represented as a (random) pricing p over lotteries, we write $\text{UTIL}(v, p)$ as the utility of the buyer; $\text{ALLOC}(v, p)$ as the allocation made to the buyer (in the form of a random subset of items or an indicator vector for that subset); and $\text{REV}(v, p)$ as the corresponding revenue of the mechanism. We extend these definitions to a value distribution \mathcal{D} by taking expectations over $v \sim \mathcal{D}$, and to a class of mechanisms \mathcal{C} by taking the maximum over all $p \in \mathcal{C}$ of the expected revenue of p . For $j \in [m]$, a subscript of j on each of these quantities indicates the contribution of item j to the corresponding quantity, where well-defined. Finally, for any valuation function v and a subset $S \subseteq [m]$ of items, we denote by $v_{|S}$ the valuation function given by $v_{|S}(T) = v(T \cap S)$ for all $T \subseteq [m]$; analogously, we denote by $\mathcal{D}_{|S}$ the value distribution that first samples a valuation $v \sim \mathcal{D}$, and then returns $v_{|S}$. In other words, $\mathcal{D}_{|S}$ captures the valuation of buyer \mathcal{D} if we can only offer them items in S .

2.1 The Ex Ante Relaxation

As discussed previously, the ex ante relaxation provides an upper bound on the revenue of an optimal mechanism by relaxing the ex post supply constraint, namely that each item should be sold at most once, to an ex ante supply constraint, namely that the expected number of copies of the item sold is at most one.

Following the approach of Chawla et al. [2023], we can define an ex ante relaxation with respect to a specific class of mechanisms. Let \mathcal{C} be a class of single-buyer Bayesian IC mechanisms, and let \mathcal{M} be a distribution over mechanisms in this class. For a buyer with value distribution \mathcal{D} , and an ex ante constraint $x \in [0, 1]^m$, we say that the distribution \mathcal{M} satisfies the ex ante constraint x over \mathcal{D} if it holds that:

$$\mathbb{E}_{M \sim \mathcal{M}} [\text{ALLOC}(\mathcal{D}, M)] \preceq x.$$

The optimal revenue achievable from a class \mathcal{C} of mechanisms given ex ante constraint x is:

$$\text{REV}_x(\mathcal{D}, \mathcal{C}) := \max_{\mathcal{M} \in \Delta^{\mathcal{C}} : \mathcal{M} \text{ satisfies } x \text{ over } \mathcal{D}} \mathbb{E}_{M \sim \mathcal{M}} [\text{REV}(\mathcal{D}, M)].$$

The optimal ex ante revenue from n buyers with joint value distribution \mathcal{D} is then defined as follows.

Definition 1 (Optimal Ex Ante Revenue under Mechanism Class). Given a class \mathcal{C} of single-buyer mechanisms and a joint value distribution \mathcal{D} , we define the *optimal ex ante revenue under \mathcal{C}* as

$$\text{EAREV}(\mathcal{D}, \mathcal{C}) := \max_{x \in \mathcal{P}} \sum_{i=1}^n \text{REV}_{x_i}(\mathcal{D}_i, \mathcal{C}),$$

where \mathcal{P} denotes the polytope of feasible ex ante allocations; that is

$$\mathcal{P} := \left\{ x = (x_1, x_2, \dots, x_n) : x_i \in \mathbb{R}_{\geq 0}^m \ \forall i \in [n] \text{ and } \sum_{i=1}^n x_{ij} \leq 1 \ \forall j \in [m] \right\}.$$

Note that the ex ante revenue $\text{EAREV}(\mathcal{D}, \mathcal{C})$ provides an upper bound on the revenue of any Bayesian IC mechanism $M = (A, \pi)$, each single-agent component of which is a \mathcal{C} mechanism; that is, for all $i \in [n]$, the distribution $(A(v_i, \cdot), \pi(v_i, \cdot))$ taken over the randomness in v_{-i} lies in $\Delta^{\mathcal{C}}$.

Chawla et al. [2023] prove the following connection between the ex ante revenue over buy many mechanisms and the ex ante item pricing revenue.

Theorem 1 ([Chawla et al., 2023]). *For any multi-buyer value distribution \mathcal{D} over m items,*

$$\text{EAREV}(\mathcal{D}, \text{BUYMANY}) \leq O(\log m) \cdot \text{EAREV}(\mathcal{D}, \text{ITEMPRICING}).$$

2.2 Our Main Results

Given the connection between the ex ante buy-many and item pricing revenues established in Theorem 1, we focus on approximating the ex ante relaxation for item pricings. Our main result is as follows:

Theorem 2. *Let \mathcal{D} be any joint distribution for n buyers and m items over subadditive valuation functions, and let SEQITEMPRICING denote the class of all Sequential Item Pricing mechanisms. Then,*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \leq O(\log m) \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

As a corollary, we obtain the following approximation to ex ante buy-many revenue.

Corollary 3. *Let \mathcal{D} be any joint distribution for n buyers and m items over subadditive valuation functions. Then we have,*

$$\text{EAREV}(\mathcal{D}, \text{BUYMANY}) \leq O(\log^2 m) \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

We show that the dependence on m in Theorem 2 is necessary for subadditive valuations: there is an $\Omega(\sqrt{\log m})$ gap between ex ante item pricing revenue and the revenue of sequential item pricings even when buyer valuations are XOS (fractionally subadditive). In fact, the lower bound also applies to the gap between the ex ante item pricing revenue and the revenue of *any* ex post feasible (random) item pricing mechanism. We prove this lower bound in Section 8.

Theorem 4. *There exists a joint value distribution \mathcal{D} for n buyers and m items over fractionally subadditive (XOS) valuation functions for which*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \geq \Omega\left(\sqrt{\log m}\right) \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

We then consider other classes of value distributions. Chawla et al. [2023] previously showed that sequential item pricings obtain a constant fraction of the ex ante item pricing revenue over the class of unit demand value distributions. We show an extension of that result to gross substitutes valuations in Appendix C. Note that the factor of 2 is tight even for $m = 1$ items.

Theorem 5. *Let \mathcal{D} be any joint distribution for n buyers and m items over gross substitutes valuation functions. Then we have,*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \leq 2 \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

Finally, in Appendix B we demonstrate upper and lower bounds for the family of general monotone valuations. Our results are summarized in Table 1.

2.3 Online Contention Resolution Schemes for Revenue

Our main goal in this paper is to approximate the ex ante revenue for a given class \mathcal{C} of mechanisms using a sequential \mathcal{C} mechanism. To this end, we define the following *Online Contention Resolution Scheme* (OCRS).

Definition 2 (OCRS for Revenue Maximization). Let \mathbb{D} be a family of single-buyer value distributions over m items and $\alpha \geq 1$. We say that \mathbb{D} admits an α -OCRS for revenue maximization under a class of mechanisms \mathcal{C} if, for any allocation constraint $x \in [0, 1]^m$, any distribution $\mathcal{D} \in \mathbb{D}$ and any random subset of items $S \subseteq [m]$, there exists a (possibly random) pricing menu $q = q(\mathcal{D}, S, x) \in \mathcal{C}$ such that:

- (a) $\mathbb{E}_S [\text{ALLOC}(\mathcal{D}_{|S}, q)] \preceq x$, and,
- (b) $\mathbb{E}_S [\text{REV}(\mathcal{D}_{|S}, q)] \geq \frac{1}{\alpha} \cdot \beta \cdot \text{REV}_x(\mathcal{D}, \mathcal{C})$, where $\beta := \min_{j \in [m]} \Pr[j \in S]$.

When \mathcal{C} is the class of item pricings, we replace (b) with a stronger condition, namely

- (b') $\mathbb{E}_S [\text{REV}(\mathcal{D}_{|S}, q)] \geq \frac{1}{\alpha} \cdot \mathbb{E}_{S,p} \left[\sum_{j \in S} p_j \text{ALLOC}_j(\mathcal{D}, p) \right]$,

where p is the (random) item pricing that achieves $\text{REV}_x(\mathcal{D}, \text{ITEMPRICING})$.

The following lemma connects OCRSes for revenue to approximations for ex ante revenue. We present its proof in Section 5.

Lemma 6. *If a family of value distributions \mathbb{D} admits an α -OCRS for revenue maximization under a class of mechanisms \mathcal{C} , then for any $\mathcal{D} \in \mathbb{D}$, we can obtain a 4α -approximation to $\text{EAREV}(\mathcal{D}, \mathcal{C})$ using a sequential \mathcal{C} mechanism.*

2.4 Establishing OCRSes through Revenue Recovery Schemes

As we saw, in order to design an OCRS for revenue maximization, we need to be able to satisfy an allocation condition as well as a revenue condition over the given set of available items. Our first main technical contribution is to show that for item pricing mechanisms we can ignore the allocation condition and instead simply require that the pricing we construct is component-wise not much smaller than the ex ante-optimal pricing. To that end, we introduce the concept of a *Revenue Recovery Scheme* (RRS).

Definition 3 (Revenue Recovery Scheme). Let \mathbb{D} be a family of single-buyer value distributions over m items and $\alpha \geq 1$. We say that \mathbb{D} admits an α -RRS if for any $\mathcal{D} \in \mathbb{D}$, any item pricing p and any subset of items $S \subseteq [m]$, there exists an item pricing $q = q(\mathcal{D}, S, p)$ such that:

- (a) $q_j \geq \frac{1}{\alpha} \cdot p_j$ for any $j \in S$, and,
- (b) $\text{REV}(\mathcal{D}_{|S}, q) \geq \frac{1}{\alpha} \cdot \sum_{j \in S} p_j \text{ALLOC}_j(\mathcal{D}, p)$.

Perhaps surprisingly, we show in Section 4, that the lower bounds on prices enforced in condition (a) are enough to guarantee an OCRS for the class of item pricing mechanisms, even though no condition on the allocation is explicitly imposed.

Lemma 7. *Any family of value distributions \mathbb{D} that admits an α -RRS also admits an $\frac{e}{e-1}\alpha$ -OCRS for revenue maximization under ITEMPRICING.*

It remains to argue that good revenue recovery schemes exist for valuation functions of interest. As a warm-up, we observe that gross substitutes valuations admit a 1-RRS. Indeed, consider the RRS that sets $q := p$. Clearly, the first scaling condition holds for $\alpha = 1$. Furthermore, we can interpret the offering of prices p over a subset S of items as being equivalent to offering the prices $q_j = p_j$ for $j \in S$ and $q_j = \infty$ for $j \notin S$. Then, by the definition of gross substitutes valuations, a buyer with value v that purchases some item $j \in S$ under p continues to purchase j under q , and the condition (b) follows. This is summarized in the following proposition. Combined with Lemma 6 this immediately implies a constant approximation to the ex ante revenue for gross substitute valuations. We show how to improve the approximation to 2 and prove Theorem 5 in Appendix C.

Proposition 8. *Gross substitutes valuations admit a 1-RRS and an $\frac{e}{e-1}$ -OCRS.*

Our main result for revenue recovery schemes applies to the much more general class of subadditive valuations and is stated through the following lemma that we prove in Section 3.

Lemma 9. *The family of subadditive valuations admits an $O(\log m + \log \Gamma)$ -RRS, where $\Gamma := p_{\max}/p_{\min}$ is the aspect ratio of the item pricing p to be revenue-recovered.*

Observe that Theorem 2 does not immediately follow by combining Lemma 6, Lemma 7, and Lemma 9 because of the dependence on the aspect ratio Γ in Lemma 9. This dependence is in fact intrinsic to the RRS definition, as there are XOS instances where it is impossible to obtain an $o(\sqrt{m})$ -RRS—a more detailed discussion can be found in Appendix D. We show how to remove this dependence for Theorem 2 in Section 6 via an end-to-end argument.

3 An RRS for Subadditive Valuations (Proof of Lemma 9)

In this section, we design a *Revenue Recovery Scheme* for *subadditive* valuations, proving Lemma 9. Fix any item pricing $p \in \mathbb{R}_{>0}^m$ and any set of items $S \subseteq [m]$. Let $\Gamma' := (\max_{j \in S} p_j) / (\min_{j \in S} p_j)$, where the minimum excludes any j such that $p_j = 0$. Although the definition of an RRS calls for a deterministic pricing q , we first show the existence of a random pricing that satisfies the conditions of the scheme.

Our algorithm samples a real $\gamma \in [\ell, h]$ from the distribution with density $f(\gamma) = 1/(\gamma \cdot \log \frac{h}{\ell})$ for $\ell = 1/2$ and $h = m\Gamma'$ and then returns a *uniform* scaling of p by γ , i.e. $q := \gamma p$. Clearly, any price is scaled down by at most $1/2$, so the scaling RRS condition holds. In order to prove Lemma 9, we will now show that for any subadditive buyer v ,

$$\text{REV}(v_{|S}, q) \geq \frac{1}{2 \log(2m\Gamma')} \cdot \sum_{j \in S} p_j \text{ALLOC}_j(v, p). \quad (1)$$

As a first step towards proving this inequality, we employ Lemma 3.1 from [Chawla et al., 2022b] that connects the revenue from any buyer with the buyer’s utility under a uniform scaling of an arbitrary pricing.

Fact 1 (Lemma 3.1 of [Chawla et al., 2022b]). *For any monotone valuation function v , any pricing p and any $0 < \ell \leq h$, if $\gamma \in [\ell, h]$ is drawn from the distribution with density $f(\gamma) = 1/(\gamma \cdot \log \frac{h}{\ell})$, then*

$$\mathbb{E}_\gamma[\text{REV}(v, \gamma p)] \geq \frac{1}{\log(h/\ell)} \cdot \left(\text{UTIL}(v, \ell p) - \text{UTIL}(v, hp) \right).$$

By applying this lemma to $v_{|S}$ and q , we immediately obtain that

$$\text{REV}(v_{|S}, q) \geq \frac{1}{\log(2m\Gamma')} \cdot \left(\text{UTIL}(v_{|S}, p/2) - \text{UTIL}(v_{|S}, m\Gamma' p) \right). \quad (2)$$

We now proceed to bound the difference in the utility terms. Observe that as a function of the scaling γ , $\text{UTIL}(v_{|S}, \gamma p) := \max_{T \subseteq S} (v(T) - \gamma p(T))$ is non-negative and non-increasing. Furthermore, as long as it remains strictly positive, it decreases at a rate of $p(T)$ for some $T \subseteq S$, which is at least $\min_{j \in S} p_j$. Thus, assuming that $\text{UTIL}(v_{|S}, hp) > 0$, we obtain that

$$\text{UTIL}(v_{|S}, \ell p) - \text{UTIL}(v_{|S}, hp) \geq (h - \ell) \cdot \min_{j \in S} p_j \geq \frac{m}{2} \cdot \max_{j \in S} p_j \geq \frac{1}{2} \cdot \sum_{j \in S} p_j \text{ALLOC}_j(v_{|S}, p)$$

under our definition of ℓ and h .

It remains to bound the utility drop if $\text{UTIL}(v_{|S}, hp) = 0$; this is precisely where the subadditivity of values is required. Let $A^* := \text{ALLOC}(v, p)$. Then, by definition we have that

$$\begin{aligned} \text{UTIL}(v_{|S}, \ell p) &:= \max_{T \subseteq S} \left(v(T) - \frac{1}{2}p(T) \right) \geq v(A^* \cap S) - \frac{1}{2}p(A^* \cap S) \\ &\geq v(A^*) - v(A^* \setminus S) - \frac{1}{2}p(A^* \cap S) \\ &= (v(A^*) - p(A^*)) - (v(A^* \setminus S) - p(A^* \setminus S)) + \frac{1}{2}p(A^* \cap S) \\ &\geq \frac{1}{2}p(A^* \cap S) = \frac{1}{2} \cdot \sum_{j \in S} p_j \cdot \text{ALLOC}_j(v, p), \end{aligned}$$

where the second line uses the subadditivity of v , the third line uses the additivity of prices and the last one uses the fact that A^* is the utility maximizing set of items for v under pricing p .

Thus, we have shown that in either case

$$\text{UTIL}(v_{|S}, \ell p) - \text{UTIL}(v_{|S}, hp) \geq \frac{1}{2} \cdot \sum_{j \in S} p_j \text{ALLOC}_j(v_{|S}, p), \quad (3)$$

and combining (3) with (2) we obtain (1) as desired. Finally, by taking the expectation of (1) over any distribution \mathcal{D} , we obtain

$$\mathbb{E}_q [\text{REV}(\mathcal{D}_{|S}, q)] \geq \frac{1}{2 \log(2m\Gamma r)} \cdot \sum_{j \in S} p_j \text{ALLOC}_j(\mathcal{D}, p).$$

and thus there must exist some *deterministic* pricing q (i.e. some *deterministic* scaling γ) for which the revenue condition holds. The proof is then completed by observing that $\Gamma' \leq \Gamma$.

4 Revenue Recovery Implies OCRS (Proof of Lemma 7)

In this section we prove Lemma 7, stating that any family of valuations that admits a α -RRS will also admit an $(e/(e-1))\alpha$ -OCRS for revenue maximization under `ITEMPRICING`. Let (\mathcal{D}, S, x) be the given instance of the OCRS. Let p denote the random pricing that achieves the revenue bound $\text{REV}_x(\mathcal{D}, \text{ITEMPRICING})$. Recall that to design an OCRS, our goal is to find a pricing over S that will recover a large fraction of $\text{REV}_x(\mathcal{D}, \text{ITEMPRICING})$ while respecting the ex ante constraint x .

To develop some intuition for this result, let us suppose that S and p are deterministic, and that the pricing p is uniform – say $p_j = 1$ for all j . Applying the RRS to (p, S) provides us with a pricing q with $q_j \geq 1/\alpha$ for all j that achieves good revenue. Since the prices q_j are comparable to the prices p_j , and the revenue of q is no more than the revenue of p , we may conclude that the total number of items allocated by q is not much larger than the total number allocated by p . In other words, the allocation constraint is satisfied *in aggregate*. In fact, we can achieve the same property over any *subset* T of S by applying the RRS to p over T to obtain pricing q_T . The aggregate allocation of each q_T is comparable to the allocation of p over the set T . We want to argue that there is a distribution over the q_T 's, corresponding to an expected allocation vector that is small in each coordinate. At the same time, we want the aggregate expected allocation of the random pricing to be comparable to that of p , so as to get back a revenue guarantee.

The key technical ingredient in our proof that enables the above properties is a novel algorithmic process called `CONVEXHULLSAMPLER` (Algorithm 1). This process takes as input a target k -dimensional vector $w \in \mathbb{R}_{\geq 0}^k$ and a collection of 2^k vectors $y^T \in \mathbb{R}_{\geq 0}^k$ corresponding to subsets $T \subseteq [k]$, with $y_j^T = 0$ for all T and $j \notin T$. The goal is to find a vector in the convex hull of the $\{y^T\}$'s that is component-wise dominated by w but at the same time comparable to w in length. We show that if each y^T is long enough, i.e. $|y^T| \geq \sum_{j \in T} w_j$, then we can always find such a vector z with $|z|/|w| \geq 1 - 1/e$.⁴

⁴Here $|z| := \sum_{j \leq k} z_j$ denotes the ℓ_1 length of the vector z .

We summarize this property of the convex hull sampler in the lemma below.

Lemma 10. *Let $w \in \mathbb{R}_{\geq 0}^k$ and $y^T \in \mathbb{R}_{\geq 0}^k$ for all $T \subseteq [k]$ satisfy (i) $y_j^T = 0$ for all T and $j \notin T$, and, (ii) $|y^T| \geq \sum_{j \in T} w_j$. Then, $\text{CONVEXHULLSAMPLER}(k, w, \{y^T\}_{T \subseteq [k]})$ returns a distribution $\{\lambda_T\}_{T \subseteq [k]}$ such that:*

- (a) $\sum_T \lambda_T y^T \preceq w$, and,
- (b) $|\sum_T \lambda_T y^T| \geq \frac{e-1}{e} \cdot |w|$.

Algorithm 1 Algorithm $\text{CONVEXHULLSAMPLER}(k, w, \{y^T\}_{T \subseteq [k]})$

- 1: Initialize variables $\tilde{w} \leftarrow w$, $Q \leftarrow [k]$, $\sigma \leftarrow 0$ and $\lambda_T \leftarrow 0$ for all $T \subseteq [k]$.
 - 2: **while** $Q \neq \emptyset$ and $\sigma < 1$ **do**
 - 3: Let τ be the largest number for which $\tau \cdot y^Q \preceq \tilde{w}$, i.e. $\tau \leftarrow \min_{j \in Q} (\tilde{w}_j / y_j^Q)$.
 - 4: Update $\lambda_Q \leftarrow \min\{\tau, 1 - \sigma\}$, $\sigma \leftarrow \sigma + \lambda_Q$ and $\tilde{w} \leftarrow \tilde{w} - \lambda_Q \cdot y^Q$.
 - 5: Update $Q \leftarrow \{i \in [k] : \tilde{w}_i > 0\}$.
 - 6: Set $\lambda_\emptyset \leftarrow 1 - \sigma$.
 - 7: Return the distribution $\{\lambda_T\}_{T \subseteq [k]}$.
-

Before we prove the lemma, let us see how it leads to a proof of Lemma 7. Our preceding discussion suggests that we should apply the Convex Hull Sampler to the allocation vectors of the pricings p and $\{q_T\}$. However, since the prices for different items can differ considerably, and in order to relate the total revenue of the pricings to the ℓ_1 lengths of the corresponding vectors, we apply the procedure to the revenue contribution of each item.

Proof of Lemma 7. We instantiate a pricing p from the distribution that generates $\text{REV}_x(\mathcal{D}, \text{ITEMPRICING})$, and instantiate a set S from its distribution. For $T \subseteq S$ let q^T be the deterministic pricing returned by the α -RRS on input (\mathcal{D}, T, p) .

Define $w_j := p_j \text{ALLOC}_j(\mathcal{D}, p)$ for $j \in S$, and $y_j^T := \alpha \cdot q_j^T \text{ALLOC}_j(\mathcal{D}_{|T}, q^T)$ for each $T \subseteq S$ and $j \in S$. We first observe that the vectors w and y^T satisfy the assumptions in Lemma 10. Assumption (i) holds by construction, and (ii) follows by noting:

$$\sum_{j \in T} y_j^T = \sum_{j \in T} \alpha \cdot q_j^T \text{ALLOC}_j(\mathcal{D}_{|T}, q^T) = \alpha \cdot \text{REV}(\mathcal{D}_{|T}, q^T) \geq \sum_{j \in T} p_j \text{ALLOC}_j(\mathcal{D}, p) = \sum_{j \in T} w_j.$$

Here the inequality follows from property (b) of the α -RRS (Definition 3).

Now, let $\lambda := \{\lambda_T\}_{T \subseteq [k]}$ be the distribution returned by CONVEXHULLSAMPLER on $(S, w, \{y^T\}_{T \subseteq S})$. Consider the following random pricing: we pick a random set T from the distribution λ and return the random pricing $\tilde{q} := q^T$.

Let us first analyze the expected allocation of this random pricing. We first note that for all T , we have $y_j^T = \alpha \cdot q_j^T \text{ALLOC}_j(\mathcal{D}_{|T}, q^T) \geq p_j \cdot \text{ALLOC}_j(\mathcal{D}_{|T}, q^T)$ from property (a) of the α -RRS (Definition 3). The expected allocation of the random pricing \tilde{q} is then given by:

$$\text{ALLOC}_j(\mathcal{D}_{|S}, \tilde{q}) = \sum_{T \subseteq S} \lambda_T \cdot \text{ALLOC}_j(\mathcal{D}_{|T}, q^T) \leq \sum_{T \subseteq S} \lambda_T \cdot \frac{y_j^T}{p_j} \leq \frac{w_j}{p_j} \leq \text{ALLOC}_j(\mathcal{D}, p),$$

where the second to last inequality follows from property (a) in Lemma 10 and the last inequality follows from the definition of w .

Next let us consider the revenue obtained by \tilde{q} . By the definition of the y^T 's, we have:

$$\text{REV}(\mathcal{D}_{|S}, \tilde{q}) = \frac{1}{\alpha} \cdot \sum_{T \subseteq S} \lambda_T |y^T| \geq \frac{1}{\alpha} \cdot \left(1 - \frac{1}{e}\right) \cdot |w| = \frac{1}{\alpha} \cdot \left(1 - \frac{1}{e}\right) \cdot \sum_{j \in S} p_j \text{ALLOC}_j(\mathcal{D}, p).$$

The lemma then follows by taking expectations for both the allocation and the revenue bounds over the randomness of S and p . \square

We conclude this section by presenting the proof of Lemma 10.

Proof of Lemma 10. We first observe by construction that $\lambda := \{\lambda_T\}_{T \subseteq [k]}$ forms a proper probability distribution. Also, it follows by construction that $\sum_T \lambda_T y^T \preceq w$. This is because we track the vector $w - \sum_T \lambda_T y^T$ in the form of vector \tilde{w} in the algorithm, and our choice of λ_Q ensures that no coordinate of this vector ever becomes negative. It remains to prove property (b).

Suppose the algorithm runs for s steps. Let $\tilde{w}_1, \tilde{w}_2, \dots, \tilde{w}_s$ and Q_1, Q_2, \dots, Q_s be the vector \tilde{w} and set Q after each step of the algorithm, and conventionally define that $\tilde{w}_0 = w$ and $Q_0 = [k]$. We first observe that

$$\left| \sum_T \lambda_T \cdot y^T \right| = \left| \sum_{i=0}^{s-1} \lambda_{Q_i} \cdot y^{Q_i} \right| = |w - \tilde{w}_s| = \sum_{i=1}^s |\tilde{w}_{i-1} - \tilde{w}_i|,$$

since $\lambda_T = 0$ for any (non-empty) T not realized in some iteration, and $w = \tilde{w}_0 \succ \tilde{w}_1 \succ \dots \succ \tilde{w}_s \succ 0$. If $\sigma < 1$ after the algorithm terminates, we necessarily have $Q_s = \emptyset$, which means that $\tilde{w}_s = 0$ and thus $|\sum_T \lambda_T \cdot y^T| = |w|$, which is clearly enough for property (b) to hold.

Thus, it remains to argue about the $\sigma = 1$ case, under which we know that $\sum_{i=0}^{s-1} \lambda_{Q_i} = 1$. For notational convenience, define the constants

$$\gamma_i = \frac{|\tilde{w}_{i-1} - \tilde{w}_i|}{|\tilde{w}_{i-1}|}, \quad i \in [s],$$

and observe that since $w = \tilde{w}_0 \succ \tilde{w}_1 \succ \dots \succ \tilde{w}_s \succ 0$, we have that $|\tilde{w}_i| = \prod_{j=1}^i (1 - \gamma_j) \cdot |w|$ and thus:

$$\left| \sum_T \lambda_T \cdot y^T \right| = \sum_{i=1}^s \gamma_i \cdot |\tilde{w}_{i-1}| = |w| \cdot \sum_{i=1}^s \left(\gamma_i \cdot \prod_{j=1}^{i-1} (1 - \gamma_j) \right) = |w| \cdot \left(1 - \prod_{i=1}^s (1 - \gamma_i) \right),$$

where the last equality follows from an expansion of the sum of products. Up next, we will argue that $\gamma_i \geq \lambda_{Q_{i-1}}$ or equivalently that $|y^{Q_{i-1}}| \geq |\tilde{w}_{i-1}|$ as $|\tilde{w}_{i-1} - \tilde{w}_i| = \lambda_{Q_{i-1}} \cdot |y^{Q_{i-1}}|$. This can be easily shown to be true for our given definitions of these vectors, as

$$|y^{Q_{i-1}}| \geq \sum_{j \in Q_{i-1}} w_j \geq \sum_{j \in Q_{i-1}} \tilde{w}_{i-1,j} = |\tilde{w}_{i-1}|,$$

where the first inequality follows from the definition of w and y^Q as well as the scaling guarantee of the RRS, the second inequality holds from $w = \tilde{w}_0 \succeq \tilde{w}_{i-1}$ and the final equality holds from the fact that all the coordinates of \tilde{w}_{i-1} not in Q_{i-1} are 0. Equipped with $\gamma_i \geq \lambda_{Q_{i-1}}$, we can finally show property (b) by observing that

$$\begin{aligned} \left| \sum_T \lambda_T \cdot y^T \right| &\geq |w| \cdot \left(1 - \prod_{i=1}^s (1 - \lambda_{Q_{i-1}}) \right) \\ &\geq |w| \cdot \left(1 - \left(\frac{\sum_{i=1}^s (1 - \lambda_{Q_{i-1}})}{s} \right)^s \right) \\ &= |w| \cdot \left(1 - \left(1 - \frac{1}{s} \right)^s \right). \end{aligned}$$

Here, the second inequality is an application of the AM-GM inequality, and the equality follows from $\sum_{i=0}^{s-1} \lambda_{Q_i} = 1$. \square

5 OCRS Approximates Ex Ante Revenue (Proof of Lemma 6)

In this section, we prove Lemma 6, which is restated for convenience. We remark that the argument we employ is standard and is used, e.g., to prove a competitive ratio of 4 for the single item prophet inequality

in [Chawla et al., 2010]. We reproduce it here for completeness. An informed reader can safely skip this section.

Lemma 6. *If a family of value distributions \mathbb{D} admits an α -OCRS for revenue maximization under a class of mechanisms \mathcal{C} , then for any $\mathcal{D} \in \mathbb{D}$, we can obtain a 4α -approximation to $\text{EAREV}(\mathcal{D}, \mathcal{C})$ using a sequential \mathcal{C} mechanism.*

Consider some $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ in \mathbb{D} and recall that

$$\text{EAREV}(\mathcal{D}, \mathcal{C}) = \sum_{i=1}^n \text{EAREV}_{x_i^*}(\mathcal{D}_i, \mathcal{C})$$

for some ex-ante constraint $\{x_i^*\}_{i=1}^n$ with $\sum_{i=1}^n x_{ij}^* \leq 1$ for all items $j \in [m]$. We will now describe a sequential \mathcal{C} mechanism; call it M . Let S_i denote the set of available (i.e. unsold) items when buyer i arrives.

1. Initially, $S_1 = [m]$.
2. When buyer $i \in [n]$ arrives:
 - (a) Present them with the pricing menu $q_i \in \mathcal{C}$ that is produced from the α -OCRS on distribution \mathcal{D}_i , subset of items S_i , and allocation constraint $\frac{1}{2}x_i^*$.
 - (b) Let $B_i \subseteq S_i$ be the set of items that the buyer purchases under this pricing; update $S_{i+1} = S_i \setminus B_i$.

Let $\text{REV}(\mathcal{D}, M)$ denote the expected revenue of the above mechanism. Note that the mechanism clearly satisfies the ex post supply constraint, as only available items are ever sold. The lemma then follows from the following sequence of steps bounding its revenue.

- (i) We have

$$\text{REV}(\mathcal{D}, M) = \sum_{i=1}^n \mathbb{E}_{S_i} [\text{REV}(\mathcal{D}_i|_{S_i}, q_i)] \geq (1/\alpha) \cdot \sum_{i=1}^n \beta_i \cdot \text{REV}_{\frac{1}{2}x_i^*}(\mathcal{D}_i, \mathcal{C}),$$

where $\beta_i = \min_{j \in [m]} \Pr[j \in S_i]$. Here the first equality is from the definition of M and the second inequality follows from the revenue condition of an α -OCRS.

- (ii) For each i and \mathcal{D}_i , we have $\text{REV}_{\frac{1}{2}x_i^*}(\mathcal{D}_i, \mathcal{C}) \geq \frac{1}{2} \cdot \text{REV}_{x_i^*}(\mathcal{D}_i, \mathcal{C})$, because one way to satisfy the ex ante constraint $\frac{1}{2}x_i^*$ is to flip a coin and use the mechanism that satisfies the ex ante constraint x_i^* with probability $1/2$ and set all prices to ∞ with probability $1/2$.
- (iii) For any item $j \in [m]$ and any buyer $i \in [n]$, it holds that $\Pr[j \in S_i] \geq \frac{1}{2}$. That is, $\beta_i \geq \frac{1}{2}$ for all i .

It remains to prove the claim in (iii), for which the following sequence of inequalities suffices:

$$\Pr[j \in S_i] = 1 - \sum_{k < i} \Pr[j \in B_k \mid j \in S_k] \geq 1 - \sum_{k < i} \frac{1}{2} x_{k,j}^* \geq \frac{1}{2}.$$

Here, the first inequality follows from the allocation condition of the OCRS (instantiated on constraints $\frac{1}{2}x_i^*$) and the second follows from the ex ante constraint. This completes the proof of Lemma 6.

6 An $O(\log m)$ -Approximation for Subadditive Valuations

Thus far we have shown through a Revenue Recovery Scheme construction that there exists a sequential item pricing mechanism that achieves an $O(\log m + \log p_{\max}/p_{\min})$ approximation with respect to $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$, whenever the valuations in \mathcal{D} are subadditive. Naturally, this dependency on the aspect ratio of the ex ante prices is not desirable, as it can potentially become very large. In this section, we show how to remove this dependency, and prove Theorem 2, which we re-state for convenience:

Theorem 2. *Let \mathcal{D} be any joint distribution for n buyers and m items over subadditive valuation functions, and let SEQITEMPRICING denote the class of all Sequential Item Pricing mechanisms. Then,*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \leq O(\log m) \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

The key observation resulting in Theorem 2 is that the $\log(p_{\max}/p_{\min})$ factor we suffer in the design of an RRS is manifested through the set of items we attempt to sell to the buyer; in other words, if we only sell items whose ex ante optimal prices lie in some window of aspect ratio $\text{poly}(m)$, then our approximation guarantee would indeed be $O(\log m)$. The problem is that by restricting what items we are allowed to sell, we will lose the potential revenue that could be collected from them. Unsurprisingly, the revenue lost from not selling *very low price items* can be neglected, as it consists only a small portion of the total revenue that we wish to approximate. However, the same argument cannot be made for the revenue lost from not selling *very high price items*, as they can potentially contribute an arbitrary large fraction of the total revenue. In fact, in Appendix D we show that this loss is necessary for some XOS value distributions – no RRS can obtain more than a $\text{poly}(m)$ fraction of the desired revenue. We circumvent this obstacle by designing a different sequential item pricing mechanism that captures the revenue contributed by the very high price items by exploiting the fact that these items are sold with polynomially small probabilities.

We will now make the above approach concrete. Consider any joint distribution \mathcal{D} and for convenience, let $\text{OBJ} = \text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$ denote the objective we wish to approximate. Let $p = (p_1, \dots, p_n)$ be the vector of (random) item pricings that achieve the optimal ex ante revenue OBJ . We will partition prices into “small”, “medium”, and “large” sets as follows. Define

$$S := \left[0, \frac{\text{OBJ}}{m^2}\right), \quad M := \left[\frac{\text{OBJ}}{m^2}, 8m^2\text{OBJ}\right], \quad L := (8m^2\text{OBJ}, +\infty).$$

Note that $\text{OBJ} = \text{S-OBJ} + \text{M-OBJ} + \text{L-OBJ}$, where

$$\text{S-OBJ} := \mathbb{E}_{v \sim \mathcal{D}, p} \left[\sum_{i=1}^n \sum_{j=1}^m p_{ij} \text{ALLOC}_j(v_i, p_i) \cdot \mathbb{1}(p_{ij} \in S) \right],$$

is the fraction of OBJ recovered via items with optimal ex ante prices in S , and M-OBJ and L-OBJ are defined analogously for M and L . We first note that S-OBJ is a negligible fraction of OBJ . Indeed,

$$\text{S-OBJ} \leq \frac{\text{OBJ}}{m^2} \cdot \mathbb{E}_{v \sim \mathcal{D}, p} \left[\sum_{i,j} \text{ALLOC}_j(v_i, p_i) \right] = \frac{\text{OBJ}}{m^2} \cdot \sum_{j=1}^m \sum_{i=1}^n \text{ALLOC}_j(\mathcal{D}_i, p_i) \leq \frac{\text{OBJ}}{m^2} \cdot \sum_{j=1}^m 1 \leq \frac{\text{OBJ}}{m},$$

with the second-to-last inequality following from the ex-ante condition. From this, observe that if we can obtain an $O(\log m)$ -approximation to M-OBJ and L-OBJ via two separate sequential item pricing mechanisms, then by flipping a fair coin at the beginning to decide which mechanism to run, we will obtain a new sequential item pricing mechanism whose expected revenue also approximates OBJ up to a factor of $O(\log m)$; this is precisely the mechanism that achieves the guarantee of Theorem 2.

As already mentioned, handling M-OBJ is easy using the machinery we established in the previous sections. When buyer i arrives, instead of attempting to sell any available item in S_i , we will only attempt to sell items $j \in S_i \cap M(p_i)$ where $p_i \in \mathbb{R}_{\geq 0}^m$ is the (sampled) ex ante price for the buyer and $M(p_i)$ denotes the set of “medium” priced items, i.e. $M(p_i) := \{j : p_{ij} \in M\}$. We can then apply the $O(\log m + \log p_{\max}/p_{\min})$ -OCRS guaranteed by the combination of Lemma 7 and Lemma 9 to $S_i \cap M(p_i)$. Using the fact that the price ratio between any two prices in M is at most $8m^4$, we will obtain a (random) pricing q_i such that $\text{ALLOC}(\mathcal{D}_{i|S_i \cap M(p_i)}, q_i) \preceq \text{ALLOC}(\mathcal{D}_i, p_i)$ and $\sum_{j \in S_i} p_{ij} \text{ALLOC}_j(\mathcal{D}_i, p_i) \leq O(\log m) \cdot \text{REV}(\mathcal{D}_{i|S_i \cap M(p_i)}, q_i)$. By taking the corresponding expectation over the sampled p_i , we obtain a $O(\log m)$ approximation to M-OBJ , as desired.

Finally, we are left with the task of obtaining a $O(\log m)$ approximation to L-OBJ ; in fact, we will show that a constant approximation is possible for all distributions over monotone valuation functions.

Lemma 11. *For any value distribution \mathcal{D} and any item pricing p , there exists an item pricing q such that $q_j \geq 2m \cdot \text{OBJ}$ for all $j \in [m]$, and*

$$\text{REV}(\mathcal{D}, q) \geq \frac{1}{4} \cdot \sum_{j=1}^m p_j \text{ALLOC}_j(\mathcal{D}, p) \cdot \mathbb{1}[p_j \in L].$$

Before proving Lemma 11, let us first argue why it directly implies a $O(1)$ approximation to L-OBJ. Consider the sequential item pricing mechanism that upon buyer i 's arrival observes the set of available (i.e. unsold) items S_i and proceeds as follows. If at least one item is unavailable, it skips the buyer, selling him nothing. Otherwise, it samples a pricing p_i from the optimal ex ante pricing for buyer i , and presents the buyer with the pricing q_i guaranteed by Lemma 11. We will prove that with constant probability this process sells no items, and therefore, obtains in expectation a constant fraction of the revenue contribution of each buyer i to L-OBJ. Formally:

Claim 12. *For each buyer i , $S_i = [m]$ with probability at least $1/2$.*

It remains to prove Lemma 11 and Claim 12.

Proof of Lemma 11. We define the item pricing q as follows.

$$q_j := \begin{cases} p_j/2 & \text{if } p_j \in L, \text{ i.e. } p_j \geq 8m^2 \cdot \text{OBJ}, \\ \max\{p_j, 2m \cdot \text{OBJ}\} & \text{otherwise.} \end{cases}$$

It is straightforward to see that $q_j \geq 2m \cdot \text{OBJ}$ for every j . To prove the lemma, it suffices to show that for any fixed buyer v it holds that

$$\text{REV}(v, q) \geq \frac{1}{4} \cdot \sum_{j=1}^m p_j \text{ALLOC}_j(v, p) \cdot \mathbb{1}(p_j \in L),$$

and then simply taking an expectation over $v \sim \mathcal{D}$ will produce the claim. Fix the buyer v and define $T := \text{ALLOC}(v, p)$ and $T' := \text{ALLOC}(v, q)$ to be utility-maximizing sets under pricings p and q respectively. We will overload notation and use L to denote the set of indices j where $p_j \in L$. Then note that the LHS of the above claimed inequality is $q(T')$ and the sum in the RHS is $p(T \cap L)$. So, we need to show that

$$q(T') \geq \frac{1}{4} \cdot p(T \cap L).$$

This clearly holds when $T \cap L = \emptyset$, so we assume that $T \cap L \neq \emptyset$. On one hand,

$$p(T) - q(T) = p(T \cap L) - q(T \cap L) + p(T \setminus L) - q(T \setminus L) \geq \frac{1}{2} \cdot p(T \cap L) - 2m^2 \cdot \text{OBJ} \geq \frac{1}{4} \cdot p(T \cap L) \quad (4)$$

where the first equality uses additivity of prices; the middle inequality uses that $p_j = 2q_j$ for all $j \in L$, while $q_j \geq 2m \cdot \text{OBJ}$ for all $j \notin L$; and the last inequality uses that $p(T \cap L) \geq 8m^2 \cdot \text{OBJ}$ as $T \cap L$ is nonempty. On the other hand,

$$p(T') - q(T') = p(T' \cap L) - q(T' \cap L) + p(T' \setminus L) - q(T' \setminus L) \leq q(T' \cap L) \leq q(T') \quad (5)$$

where we used $q_j \geq p_j$ for $j \notin L$. Finally, as $v(T) - p(T) \geq v(T') - p(T')$ and $v(T') - q(T') \geq v(T) - q(T)$ by the definition of T and T' , we have $q(T) - q(T') \geq v(T) - v(T') \geq p(T) - p(T')$, or

$$p(T') - q(T') \geq p(T) - q(T). \quad (6)$$

The lemma then follows from combining Equations (4) to (6). \square

Proof of Claim 12. It suffices to show that $\sum_{i' < i} \sum_{j=1}^m \text{ALLOC}_j(\mathcal{D}_{i'}, q_{i'}) \leq \frac{1}{2}$, from which our claim follows from a union bound. Notice that here $q_{i'}$ is a random variable that depends on the pricing $p_{i'}$ that our algorithm draws.

We first note that $\text{ALLOC}_j(\mathcal{D}_{i'}, q_{i'}) \leq \frac{1}{2m}$ for every buyer i' ⁵; otherwise, as $q_{i',j} \geq 2m \cdot \text{OBJ}$, by simply using q for the buyer i' with valuation $\mathcal{D}_{i'}$ (and not selling anything to any other buyer) we would get an expected revenue that is strictly larger than $\text{OBJ} := \text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$ which is obviously a contradiction by the optimality of the ex-ante solution. This means that for every buyer i' , we have

$$\sum_{j=1}^m \text{ALLOC}_j(\mathcal{D}_{i'}, q_{i'}) \leq \frac{1}{2}.$$

We now prove the claim via contradiction. Suppose the claim is not true for some $i \in [n]$. Then there must exist a first index $k \leq i$ where it becomes not true, namely,

$$\sum_{i' < k} \sum_{j=1}^m \text{ALLOC}_j(\mathcal{D}_{i'}, q_{i'}) \leq \frac{1}{2} \quad \text{and} \quad \sum_{i' \leq k} \sum_{j=1}^m \text{ALLOC}_j(\mathcal{D}_{i'}, q_{i'}) > \frac{1}{2}.$$

Note that as $\sum_{j=1}^m \text{ALLOC}_j(\mathcal{D}_k, q_k) \leq \frac{1}{2}$, we have

$$\sum_{i' \leq k} \sum_{j=1}^m \text{ALLOC}_j(\mathcal{D}_{i'}, q_{i'}) \leq 1.$$

Therefore, the pricing vector q' which is equal to q on the first k buyers and to infinity on the remaining buyers satisfies the ex-ante constraint, and its ex-ante revenue will be at least

$$\mathbb{E} \left[\sum_{i' \leq k} \sum_{j=1}^m q_j \cdot \text{ALLOC}_j(\mathcal{D}_{i'}, q_{i'}) \right] \geq \frac{1}{2} \cdot 2m \cdot \text{OBJ} \geq m \text{OBJ},$$

which contradicts ex-ante optimality. \square

7 Computational Considerations

In this section, we summarize Sections 3 through 6 to present a top-down overview for the existence of the sequential item pricing mechanism that achieves the guarantees of Theorem 2, namely the $O(\log m)$ -approximation to the ex ante revenue under subadditive valuations, as well as comment on the computational aspects of our approach.

Computational Aspects. Recall that our objective is to approximate

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) = \sum_{i=1}^n \text{EAREV}_{x_i^*}(\mathcal{D}_i, \text{ITEMPRICING})$$

under some (optimal) ex-ante allocation constraint $\{x_i^*\}_{i=1}^n$. Our approach crucially relies on access to this optimal solution, as well as the distributions (over item pricings) p_i^* that achieve $\text{EAREV}_{x_i^*}(\mathcal{D}_i, \text{ITEMPRICING})$ for each buyer $i \in [n]$. We comment that finding the revenue-optimal item pricing is hard even in the absence of an ex ante constraint [Chalermsook et al., 2013], but approximations may exist in special cases (e.g., [Chawla et al., 2022a]).

⁵Observe that from this, if the ex ante allocations are at least $1/(2m)$ then we can cast this algorithm as a $O(1)$ -OCRS w.r.t. L-OBJ. Combined with the $O(\log m)$ -OCRS w.r.t. M-OBJ, this justifies the claim of Informal Theorem 1.

Our approach also requires access to an oracle that, given distribution \mathcal{D}_i and pricing q , can compute the expected allocation $\text{ALLOC}(\mathcal{D}_i, q)$ and expected revenue $\text{REV}(\mathcal{D}_i, q)$. Note that if the supports of the buyers' value distributions are small and we have access to a demand oracle over subadditive values, we can compute these quantities.

We remark that assuming access to x_i^* , p_i^* , and the demand oracle as discussed above, the remainder of our algorithm is *fully constructive*, i.e. we can efficiently construct the Sequential Item Pricing achieving our claimed bounds in time polynomial in n , m , and the support sizes of the value distributions.

Construction Overview. As described in Section 6, our construction begins by flipping a fair coin in order to decide whether it will focus on *high-price items* or *medium-price items* throughout the execution.

If the coin lands on high-price items, we proceed to address the buyers sequentially, until we sell at least one item to a buyer; once this happens, we completely ignore the rest of the buyers. To address buyer i , we sample an item pricing p_i from distribution p_i^* and present them with the pricing $q_i = q_i(p_i)$ that we define in Lemma 11. Assuming access to the value of the (optimal) ex ante revenue, this construction is straightforward and thus the entire sequential item pricing is constructed in $O(nm)$ time. In fact, since q_i doesn't depend on the set of available items, we can pre-compute these prices in advance, i.e. our sequential item pricing is oblivious to the order in which the buyers arrive, as well as the randomness of their valuations.

If the coin lands on medium-price items, then whenever buyer i arrives we observe the set of available items S_i and offer them a realization of the (random) item pricing $q_i = \text{OCRS}(\mathcal{D}_i, S_i, \frac{1}{2}x_i^*)$. Note that this sequential item pricing is adaptive both to the order in which the buyers arrive, and to the realizations of their values, as both affect the set of available items S_i . Furthermore, to avoid having to sample from the optimal pricing distribution that respects $\frac{1}{2}x_i^*$, we can instead present them with the pricing $q_i = \text{OCRS}(\mathcal{D}_i, S_i, x_i^*)$ with 1/2 probability (and entirely skip the buyer with the remaining 1/2 probability); it is easy to see that the proof of Section 5 also holds for this alternative.

In order to construct $q_i = \text{OCRS}(\mathcal{D}_i, S_i, x_i^*)$, we sample an item pricing p_i from distribution p_i^* . We then compute the vector $w_j = p_{ij} \text{ALLOC}_j(\mathcal{D}_i, p_i)$ for $j \in S_i$ as in Section 4, and then run the algorithm `CONVEXHULLSAMPLER`. An important observation is that instead of pre-computing the y^T vectors for all $T \subseteq S_i$, we can compute them on demand; then, since `CONVEXHULLSAMPLER` clearly runs for at most m steps, we only need to compute m of these vectors.

To compute a y^T vector, we need to find a pricing q^T satisfying the revenue guarantee, and then compute its revenue. From Section 3, we know that picking the best scaling $\gamma \in [\ell, h]$ and setting $q = \gamma \cdot p$ suffices. In fact, because of the linear dependence of the price of any particular set of items on γ , at the loss of a factor of 2 it suffices to choose γ as the best power of 2 in the range $[\ell, h]$.

Once `CONVEXHULLSAMPLER` completes its execution and returns the distribution λ_T over the subsets $T \subseteq S_i$, we simply sample such a set T from it and return $q_i = q^T$.

8 Lower Bounds for XOS Valuations (Proof of Theorem 4)

In this section, we prove that the gap between sequential item pricings and ex ante item pricings can be as large as $\Omega(\sqrt{\log m})$ under XOS valuations, as stated in Theorem 4, which we restate for convenience:

Theorem 4. *There exists a joint value distribution \mathcal{D} for n buyers and m items over fractionally subadditive (XOS) valuation functions for which*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \geq \Omega\left(\sqrt{\log m}\right) \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

In fact, a stronger separation holds. Let `EXPOSTITEMPRICING` be the class of multi-buyer mechanisms which take all buyers' valuations v_i as input, and offer each buyer a randomized item pricing $M_{i, v_{-i}}$ over the items which depends only on \mathcal{D}_i and the other buyers' valuations $v_{-i} := \{v_{i'} : i' \neq i\}$. Provided that the allocations under the $M_{i, v_{-i}}$ are feasible, these are feasible and truthful multi-buyer mechanisms. Furthermore, this generalizes `SEQITEMPRICING`, since any randomized sequential item pricing mechanism

can be simulated by a mechanism in EXPOSTITEMPRICING. We will show that the separation holds for this class of mechanisms as well.

Theorem 13. *There exists a joint value distribution \mathcal{D} for m items and $n = \sqrt{m}$ buyers for which*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \geq \Omega\left(\sqrt{\log m}\right) \cdot \text{REV}(\mathcal{D}, \text{EXPOSTITEMPRICING}).$$

Theorem 4 then follows as a corollary of Theorem 13, since every mechanism $M \in \text{SEQITEMPRICING}$ may be simulated by a mechanism in EXPOSTITEMPRICING.

The rest of this section is devoted to proving Theorem 13. We begin by specifying our instance. We first define a set of parameters that will be useful towards stating our claims; let $k := \sqrt{m}$, and $t := \sqrt{\log k}$, taking logarithms to be base two. By incurring at most a constant factor, we can assume that all of these parameters are integer-valued, and also that t is even.

The buyers. Conveniently, the n buyers $v_i \sim \mathcal{D}_i$ are all (independently) identically distributed, with valuations sampled according to the following process:

1. A set of items $A \subseteq [m]$ with $|A| = k$ is sampled uniformly at random.
2. An integer $h \in \{1, 2, \dots, \frac{1}{2} \log k\}$ is drawn uniformly at random. We denote $\ell := 2^h$.
3. The valuation of the buyer is then realized as an XOS function, which is given by the maximum over the following additive valuations:
 - (a) For the set A , we define an additive valuation v_A such that $v_A(j) = 1 + t$ for items $j \in A$ and $v_A(j) = 0$ for items $j \notin A$.
 - (b) For *every* set B of cardinality $|B| = t \cdot \ell$, we define an additive valuation v_B such that $v_B(j) = 1 + \frac{k}{\ell}$ for items $j \in B$ and $v_B(j) = 0$ for items $j \notin B$.

Observe that for our selected parameters $t \cdot \ell \leq m$ always, and all parameters are integer-valued, so this is a well-defined XOS instance. Let A_i be the A -set for buyer i .

Informally, this valuation ensures that each buyer i is nearly indifferent between the (single) high-revenue set A_i on the one hand, and one of the (many) low-revenue sets B on the other. For any sequential item pricing mechanism, as the set of available items shrinks, A_i becomes less valuable to i , while the most valuable low-revenue set B does not. As a result, it quickly becomes unlikely that the available fraction of A_i is large enough for i to choose it, resulting in reduced revenue.

Ex-ante revenue. To demonstrate a gap, we need a lower bound on $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$ for the above instance \mathcal{D} . The key observation is that ex ante, all buyers can be allocated their A -sets. Consider the item pricing $p_{ij} = 1$ for all buyers $i \in [n]$ and items $j \in [m]$, meaning that each buyer is presented with the exact same set of uniform prices over the items. For this pricing, the utility of the A set and any B sets are precisely

$$\begin{aligned} \text{UTIL}_i(A) &= v_i(A) - p_i(A) = k(1 + t - 1) = kt, \quad \text{and,} \\ \text{UTIL}_i(B) &= v_i(B) - p_i(B) = t\ell \left(1 + \frac{k}{\ell} - 1\right) = kt. \end{aligned}$$

By breaking ties for A (or by slightly increasing A 's valuation by $\epsilon > 0$) we can assume the buyer will always buy set A . Since A is a uniformly random subset of $[m]$ with cardinality k , each item will therefore be sold to the buyer with probability $k/m = m^{-1/2}$. Since we have $n = \sqrt{m}$ identical buyers in total, this means that the fractional total allocation of each item is precisely 1, and thus the ex ante constraint is satisfied for our proposed pricing p , implying that

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \geq \sum_{i=1}^n \text{REV}(\mathcal{D}_i, p_i).$$

Since we have already showed that under this p_i buyer i always purchases A_i , the per-buyer revenue is simply $\text{REV}(\mathcal{D}_i, p_i) = |A_i| = k$ and so for our choice of parameters $n = k = \sqrt{m}$,

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \geq n \cdot k = m. \quad (7)$$

Revenue and allocation for single-buyer item pricings. The following are crucial to analyzing this instance. They establish conditions under which a buyer can be sold items according to their A -type versus B -type additive valuations, and upper bounds on the expected revenue in each case. Omitted proofs appear in Appendix E. We begin with claims that hold for all valuations $v \sim \mathcal{D}_i$ from the above distribution, all item pricings p , and all subsets of available items $S \subseteq [m]$. Let $\mathcal{E}^A(v_i, p)$ denote the event that buyer v_i facing prices p chooses an allocation where their valuation is supported by v_A ; define $\mathcal{E}^B(v_i, p)$ analogously when their valuation is supported by any of the v_B .

Lemma 14. *For any v_i, p such that $\mathcal{E}^A(v_i, p)$ holds, buyer i purchases at least $\frac{kt}{t+1}$ items.*

Lemma 15. *For some $c_1 \in \mathbb{R}_{\geq 0}$ and sufficiently large k , for all p , $\mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^A(v_i, p)] \leq c_1 \cdot k$.*

Lemma 16. *For some $c_2 \in \mathbb{R}_{\geq 0}$ and sufficiently large k , for all p , $\mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p)] \leq c_2 \cdot t \cdot \frac{k}{\log k}$.*

These will enable us to reason about item-pricing mechanisms in multi-buyer settings for this instance.

Revenue and allocation for multi-buyer item pricings. Lemma 14 implies that no buyer facing an item pricing will choose an A -type allocation unless they receive a large proportion of their A -set A_i . We will now argue that with high probability only a few buyers can receive large proportions of their A -sets (and therefore choose a A -type allocations), even if the mechanism is given access to A_1, \dots, A_n upfront.

We will make use of the abstraction of an *assignment*, which is an arbitrary mapping of items to buyers.

Definition 4. Let $\sigma : [m] \rightarrow [n] \cup \{\perp\}$ be an *assignment* of items to buyers.

Note that every multi-buyer mechanism—item-pricing or otherwise—induces a distribution over assignments, where $\sigma^{-1}(\perp)$ are the items not allocated to any buyer. We will call an allocation *feasible* for a buyer if it allocates them at least a $(1 - c)$ -proportion of their A_i :

Definition 5. A buyer $i \in [n]$ is $(1 - c)$ -feasible under σ if $|A_i \cap \sigma^{-1}(i)| \geq (1 - c)n$. For $I \subseteq [n]$, I is $(1 - c)$ -feasible under σ if all $i \in I$ are $(1 - c)$ -feasible under σ .

Using these definitions we present the following combinatorial lemma, which together with Lemma 14 implies that for any multi-buyer mechanism which presents buyers with item pricings, the number of buyers with A -type allocations is small with high probability:

Lemma 17. *Let $I_{\max}(v)$ be the largest $I \subseteq [n]$ for which I is $(1 - \frac{1}{t+1})$ -feasible under σ , for any assignment σ . Then for some c_3 and sufficiently large k , $\mathbb{E}_v [|I_{\max}(v)|] \leq c_3 \cdot \frac{n}{t}$.*

We are now ready to prove Theorem 13.

Proof of Theorem 13. The proof is intuitive: with high probability, no mechanism can allocate A -sets to more than a $(1/t)$ -proportion of the buyers and get revenue k from each, and since it presents buyers with item pricings, its revenue on the remaining buyers is at most $tk/\log k$ each. Then, $t = \sqrt{\log k}$ balances these terms, resulting in the gap.

Formally, fix an ex post randomized item pricing mechanism $M \in \text{EXPOSTITEMPRICING}$. Then M is a collection of randomized item pricings $M_{i,v_{-i}}$, and the expected revenue is

$$\begin{aligned} \text{REV}(\mathcal{D}, M) &= \mathbb{E}_v \left[\sum_i \left(\mathbb{E}_{p_i \sim M_{i,v_{-i}}} [\text{REV}(v_i, p_i)] \right) \right] \\ &= \mathbb{E}_v \left[\sum_i \left(\mathbb{E}_{p_i \sim M_{i,v_{-i}}} \left[\mathbb{E}_{v_i} [\text{REV}(v_i, p_i)] \right] \right) \right], \end{aligned}$$

because $M_{i,v_{-i}}$ is independent of v_i . Then conditioning on the events $\mathcal{E}^A := \mathcal{E}^A(v_i, p_i)$ and $\mathcal{E}^B := \mathcal{E}^B(v_i, p_i)$,

$$\begin{aligned} &= \mathbb{E}_v \left[\sum_i \left(\mathbb{E}_{p_i \sim M_{i,v_{-i}}} \left[\mathbb{E}_{v_i} [\text{REV}(v_i, p_i) \mid \mathcal{E}^A] \cdot \Pr_{v_i} [\mathcal{E}^A] + \mathbb{E}_{v_i} [\text{REV}(v_i, p_i) \mid \mathcal{E}^B] \cdot \Pr_{v_i} [\mathcal{E}^B] \right] \right) \right] \\ &\leq \mathbb{E}_v \left[\sum_i \left(\mathbb{E}_{p_i \sim M_{i,v_{-i}}} \left[c_1 \cdot k \cdot \Pr_{v_i} [\mathcal{E}^A(v_i, p_i)] + c_2 \cdot \frac{tk}{\log k} \right] \right) \right] \\ &= c_2 \cdot \frac{ntk}{\log k} + c_1 \cdot k \cdot \mathbb{E}_v \left[\sum_i \left(\mathbb{E}_{p_i \sim M_{i,v_{-i}}} \left[\Pr_{v_i} [\mathcal{E}^A(v_i, p_i)] \right] \right) \right], \end{aligned}$$

for some constants c_1, c_2 , by Lemmas 15 and 16. Next, observe that for fixed v, p the allocations $\text{ALLOC}(v_i, p_i)$ form an assignment σ , since M is a feasible mechanism. Let $I(v, p)$ be the set of buyers for which this assignment is $(1 - \frac{1}{t+1})$ -feasible. Then Lemma 14 implies that $\mathcal{E}^A(v_i, p_i) \subseteq \{i \in I(v, p)\}$, so

$$\begin{aligned} &\leq c_2 \cdot \frac{ntk}{\log k} + c_1 \cdot k \cdot \mathbb{E}_v \left[\sum_i \left(\mathbb{E}_{p_i \sim M_{i,v_{-i}}} \left[\Pr_{v_i} [i \in I(v, p)] \right] \right) \right] \\ &= c_2 \cdot \frac{ntk}{\log k} + c_1 \cdot k \cdot \mathbb{E}_v \left[\mathbb{E}_{p \sim M} [|I(v, p)|] \right]. \end{aligned}$$

Next let $I_{\max}(v)$ be a maximum-size $(1 - \frac{1}{t+1})$ -feasible subset of the buyers v over *any* assignment $\sigma : [m] \rightarrow [n] \cup \{\perp\}$, and note that $|I(v, p)| \leq |I_{\max}(v)|$. Then by Lemma 17 we have that $\mathbb{E}_v [|I_{\max}|] \leq c_3 n/t$ for some constant c_3 . Therefore

$$\begin{aligned} \text{REV}(\mathcal{D}, M) &\leq c_2 \cdot \frac{ntk}{\log k} + c_1 \cdot k \cdot \mathbb{E}_v [|I_{\max}(v)|] \\ &\leq c_2 \cdot \frac{ntk}{\log k} + c_1 \cdot k \cdot \frac{c_3 \cdot n}{t}. \end{aligned}$$

By our choice of parameters $n = k = \sqrt{m}$ and $t = \sqrt{\log k}$ this implies that $\text{REV}(\mathcal{D}, M) = O(m \log^{-1/2} m)$, while by (7) the ex ante item pricing revenue for \mathcal{D} is at least m . This holds for all $M \in \text{EXPOSTITEMPRICING}$, completing the proof. \square

9 Related Work

Multi-parameter Revenue Maximization. As mentioned previously, multi-parameter revenue maximization is a notoriously hard problem. Briest et al. [2015] and Hart and Nisan [2019] showed that the optimal revenue is inapproximable within any finite factor by simple mechanisms (such as item pricings) even with just two items and one unit-demand or additive buyer. Much of the literature on simplicity versus optimality for revenue maximization accordingly makes strong assumptions on the buyers' value distributions, such as requiring the values to be independent across items. Under these assumptions, mechanisms

such as sequential item pricing, grand bundle pricing and two part tariffs obtain constant factor approximations to the optimal revenue. See, e.g., [Chawla et al., 2015, Babaioff et al., 2014, Rubinstein and Weinberg, 2018, Cai and Zhao, 2017, Cai et al., 2019].

There are several connections between this literature and our work. Chawla et al. [2010] showed that under the item independence assumption, the revenue maximization problem for *unit demand* buyers under various feasibility constraints can be reduced to a single-parameter prophet inequality. They were the first to use an ex ante relaxation in multi-buyer settings, a technique that was later formalized by Alaei [2014]. Chawla and Miller [2016] extend a similar approach based on the ex ante relaxation and prophet inequalities to matroid rank values, still with an independence-across-items assumption.

It is worthwhile contrasting our main results with those of Cai and Zhao [2017]. They consider buyers with subadditive values and bound the gap between the revenue of a sequential two-part tariff mechanism and the ex ante relaxation. There are several similarities and differences from our work: (1) Like the papers referenced above, Cai and Zhao require values to be independent across items, a significant restriction. Our results do not require any distributional assumptions. (2) Under the independence assumption however, they are able to compete against the ex ante optimal revenue, where our approximations only apply to the ex ante buy-many revenue, a weaker benchmark. (3) Cai and Zhao’s mechanism is a two-part tariff, meaning that every buyer needs to pay an *entry fee* to enter the mechanism and is then offered an item pricing. Two-part tariffs are necessary to obtain constant factor approximations to the optimal revenue. (4) Finally, Cai and Zhao obtain a constant factor gap for XOS valuations and an $O(\log m)$ gap for subadditive valuations, whereas we obtain an $O(\log^2 m)$ gap for both. The $O(\log m)$ factor for subadditive values was subsequently improved to an $O(\log \log m)$ factor by Dütting et al. [2020b] via the construction of a new prophet inequality for welfare maximization over subadditive values.

Buy-Many Mechanisms. Another line of work on simple versus optimal multi-parameter mechanisms considers approximating a weaker benchmark, namely *buy many* mechanisms, rather than introducing assumptions on the value distributions. This line of work was introduced by Briest et al. [2015] for unit-demand buyers and extended to general buyers by Chawla et al. [2022b]. These papers consider the single-buyer mechanism design problem and bound the gap between the revenue of item pricing and the optimal buy many mechanism. This gap is shown to match the best possible gap achievable by *any* mechanism with subexponential description complexity. Chawla et al. [2023] were the first to consider buy-many mechanisms in multi-buyer settings. They argue that any reasonable extensions of buy-many mechanisms to multiple buyers are upper bounded by ex ante buy-many revenue. As mentioned previously, this work bounds the gap between sequential item pricing revenue and the ex ante buy-many revenue, a result we extend to subadditive buyers.

Prophet inequalities and contention resolution. Prophet inequalities were first introduced to mechanism design by Hajiaghayi et al. [2007]. Chawla et al. [2010] showed how to apply them to revenue maximization settings and motivated the design of prophet inequalities for more general settings. Further work on prophet inequalities has considered variants with different feasibility constraints, arrival models, limited information settings, buyer correlations, etc. We refer the reader to the survey by Lucier [2017] for an overview of this work.

As mentioned previously, single-parameter prophet inequalities have strong connections with Online Contention Resolution Schemes. Contention Resolution Schemes were first defined by Vondrák et al. [2011] in the context of offline maximization of set functions under various feasibility constraints as an abstraction for rounding the fractional relaxations of these objectives. Subsequently Feldman et al. [2016] defined *Online* CRSes as analogous online rounding primitives for a range of online adversaries. Feldman et al. observed that good OCRSes imply good prophet inequalities for combinatorial constraints and linear objectives. Using duality, Lee and Singla [2018] showed that the converse also holds – good prophet inequalities imply good OCRSes.

Combinatorial prophet inequalities. Prophet inequalities have also been studied extensively in the multi-item setting for the social welfare objective where buyers have combinatorial valuations over the items. This line of work was introduced by Feldman et al. [2015]. Inspired by the matroid prophet inequality of Kleinberg and Weinberg [2012], Feldman et al. introduced the approach of using *balanced prices* to design prophet inequalities. This has become the dominant technique for designing combinatorial prophet inequalities (see, e.g., [Dütting et al., 2020a,b,a]), although some constructions use duals or other approaches for constructing prices [Zhang, 2022, Correa and Cristi, 2023]. We remark that all of these approaches are quite different from OCRSes in that they do not necessarily preserve the probabilities of allocation of the ex ante relaxation. Finally, we note that for the social welfare objective over subadditive values, Correa and Cristi [2023] recently developed a constant factor prophet inequality.

10 Conclusion and Open Problems

Our work develops polylogarithmic upper and lower bounds on the ratio between the revenue obtained by sequential item pricings and the ex ante item pricing (or buy many) revenue. Although the $O(\log m)$ gap between the sequential item pricing revenue and the ex ante buy-many revenue is tight for unit-demand valuations, it is not clear whether an additional $\log m$ factor is necessary to lose for XOS or subadditive values.

Another important direction is to make our result algorithmic. Constructing our item pricing requires access to the optimal solution to the ex ante item pricing relaxation. While optimizing item pricing revenue is hard even in the absence of ex ante constraints, *approximations* may be possible under suitable assumptions. Another possible avenue is to first construct an approximately optimal ex ante buy-many mechanism and then convert it into an approximately optimal ex ante item pricing. The complexity of optimizing revenue over buy-many mechanisms is not well understood. Furthermore, we note that the upper bound on the gap between the ex ante item pricing and buy many revenues (due to [Chawla et al., 2023]) is non constructive.

References

- Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.
- Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 21–30. IEEE, 2014.
- Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing lotteries. *Journal of Economic Theory*, 156:144–174, 2015.
- Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, page 170–183. Association for Computing Machinery, 2017.
- Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality-based unified approach to bayesian mechanism design. *SIAM Journal on Computing*, 50(3):STOC16–160, 2019.
- Parinya Chalermsook, Bundit Laekhanukit, and Danupon Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 370–379. IEEE, 2013.
- Shuchi Chawla and J Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 579–596, 2016.
- Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.
- Shuchi Chawla, David Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
- Shuchi Chawla, Rojin Rezvan, Yifeng Teng, and Christos Tzamos. Pricing ordered items. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, page 722–735. Association for Computing Machinery, 2022a.
- Shuchi Chawla, Yifeng Teng, and Christos Tzamos. Buy-many mechanisms are not much better than item pricing. *Games Econ. Behav.*, 134:104–116, 2022b.
- Shuchi Chawla, Rojin Rezvan, Yifeng Teng, and Christos Tzamos. Buy-many mechanisms for many unit-demand buyers. In *Web and Internet Economics - 19th International Conference, WINE 2023*, volume 14413, pages 21–38. Springer, 2023.
- José Correa and Andrés Cristi. A constant factor prophet inequality for online combinatorial auctions. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 686–697. Association for Computing Machinery, 2023.
- Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3):540–582, 2020a.
- Paul Dütting, Thomas Kesselheim, and Brendan Lucier. An $o(\log \log m)$ prophet inequality for subadditive combinatorial auctions. *SIAM Journal on Computing*, pages 239–275, 2020b.
- Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’15, page 123–135. Society for Industrial and Applied Mathematics, 2015.

- Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 1014–1033. SIAM, 2016.
- Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 1, AAAI’07*, page 58–65. AAAI Press, 2007.
- Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory*, 183:991–1029, 2019.
- Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.
- Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing, STOC ’12*, page 123–136. Association for Computing Machinery, 2012.
- Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms, ESA 2018*, volume 112 of *LIPIcs*, pages 57:1–57:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- Brendan Lucier. An economic view of prophet inequalities. *SIGecom Exch.*, 16(1):24–47, sep 2017.
- Aviad Rubinstein and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):1–25, 2018.
- Jan Vondrák, Chandra Chekuri, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011*, pages 783–792. ACM, 2011.
- Hanrui Zhang. Improved prophet inequalities for combinatorial welfare maximization with (approximately) subadditive agents. *Journal of Computer and System Sciences*, 123:143–156, 2022.

A Mechanism Design Basics

In this section we cover the necessary background on mechanism design that is required for this paper. We consider a setting with m items and n buyers with combinatorial valuations over the items, with the objective of maximizing the total revenue of the seller. We begin by establishing some notation.

Buyers and valuations. A *buyer* is modeled through a *valuation function* $v : 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$ that is drawn from a known distribution \mathcal{D} and assigns a non-negative value to each subset of the m items. Some of our results reference special classes of valuations that we define below.

- **Unit-Demand:** A valuation v is *unit-demand* if $v(S) = \max_{j \in S} v(\{j\})$ for all $S \subseteq [m]$. For convenience, we write $v(\{j\})$ as simply v_j .
- **Gross Substitutes:** Let p and p' be two item pricings with $p \preceq p'$ and let $S = \{j : p_j = p'_j\}$ denote the set of items where the two pricings charge equal prices. We say that a value function v satisfies *gross substitutes* if for all such pricings p and p' , $\text{ALLOC}(v, p) \cap S \subseteq \text{ALLOC}(v, p') \cap S$. In other words, under the new pricing p' the buyer continues to purchase those items in $\text{ALLOC}(v, p)$ whose prices did not increase.
- **Fractionally Subadditive (XOS):** A valuation v is *fractionally subadditive* if there exists a set \mathcal{A} of m -dimensional vectors $a \in \mathbb{R}_{\geq 0}^m$, such that $v(S) = \max_{a \in \mathcal{A}} \sum_{j \in S} a_j$ for all $S \subseteq [m]$.
- **Subadditive:** A valuation v is *subadditive* if $v(S) + v(T) \geq v(S \cup T)$ for all $S, T \subseteq [m]$.
- **Monotone:** A valuation v is *monotone* if $v(S) \leq v(T)$ for all $S \subseteq T \subseteq [m]$.

Observe that we have Unit-Demand \subsetneq Gross Substitutes \subsetneq XOS \subsetneq Subadditive \subsetneq Monotone.

Finally, for any valuation function v and a subset $S \subseteq [m]$ of items, we denote by $v_{|S}$ the valuation function given by $v_{|S}(T) = v(T \cap S)$ for all $T \subseteq [m]$; analogously, we denote by $\mathcal{D}_{|S}$ the value distribution that first samples a valuation $v \sim \mathcal{D}$, and then returns $v_{|S}$. In other words, $\mathcal{D}_{|S}$ captures the valuation of buyer \mathcal{D} if we can only offer them items in S .

Bayesian Incentive Compatible Mechanisms. In the multi-buyer setting, we will use \mathcal{D}_i to denote the value distribution of buyer i and $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ to denote the joint value distribution of all n buyers. Note that each buyer draws his value independently from his respective distribution. For $i \in [n]$, v_i denotes the value function of buyer i , v denotes the joint value function of all n buyers, and v_{-i} the joint value function of all buyers other than i .

A multi-buyer mechanism maps the joint value function v to an outcome $A(v)$ and a payment $\pi(v)$ where $A_i(v)$ denotes the set of items allocated to buyer i and $\pi_i(v)$ denotes the corresponding payment made by buyer i . A mechanism (A, π) is *incentive compatible* if no buyer can benefit from misreporting his value regardless of others' values. In other words, if for all i , v_i , v'_i , v_{-i} , it holds that

$$v_i(A_i(v_i, v_{-i})) - \pi_i(v_i, v_{-i}) \geq v_i(A_i(v'_i, v_{-i})) - \pi_i(v'_i, v_{-i})$$

A *Bayesian incentive compatible* (BIC) mechanism (A, π) is a pair for which the above inequality holds for all i , v_i , and v'_i in expectation over $v_{-i} \sim \mathcal{D}_{-i}$. The revenue of a multi-buyer BIC mechanism $M = (A, \pi)$ is given by $\text{REV}(\mathcal{D}, M) := \mathbb{E}_{v \sim \mathcal{D}}[\sum_{i \in [n]} \pi_i(v)]$. We use $\text{REV}(\mathcal{D}) := \max_{\text{BIC mechanism } M} \text{REV}(\mathcal{D}, M)$ to denote the optimal revenue achievable by any BIC mechanism over the joint distribution \mathcal{D} .

Single-Buyer Mechanisms as Pricings. From the viewpoint of a single buyer i , any given mechanism $M = (A, \pi)$ generates as a function of the buyer's valuation v_i a distribution over outcomes $(A(v_i, \cdot), \pi(v_i, \cdot))$. In the following discussion of single parameter mechanisms, where clear from context, we will drop the subscript i to simplify notation. By the *taxation principle*, we can interpret this single buyer mechanism as

a (possibly random) price menu $p : \Delta^{[m]} \rightarrow \mathbb{R}_+$ that maps each distribution⁶ over subsets of $[m]$, a.k.a. a lottery, to a price the buyer must pay to receive it. We will assume without loss of generality that $p(\emptyset) = 0$ with probability 1. Under this price menu, the buyer's utility for a (random) set or lottery $S \subseteq [m]$ of items is given by

$$\text{UTIL}(v, p, S) := \mathbb{E}_S \left[v(S) - p(S) \right].$$

The buyer purchases their utility maximizing random set, and the seller collects the associated revenue:

$$\text{ALLOC}(v, p) := \arg \max_{S \in \Delta^{[m]}} \text{UTIL}(v, p, S),$$

$$\text{REV}(v, p) := p \left(\text{ALLOC}(v, p) \right).$$

In the case of ties, we assume that the buyer purchases any set of maximal price from among the ones that maximize their utility. For simplicity we will overload notation and also let $\text{ALLOC}(v, p)$ denote the indicator vector of the set purchased by the buyer, and further abuse notation by calling the average such indicator vector the *expected allocation*, and denoting it and the expected revenue of the mechanism p by

$$\begin{aligned} \text{ALLOC}(\mathcal{D}, p) &:= \mathbb{E}_{p, v \sim \mathcal{D}} [\text{ALLOC}(v, p)], \\ \text{REV}(\mathcal{D}, p) &:= \mathbb{E}_{p, v \sim \mathcal{D}} [\text{REV}(v, p)]. \end{aligned}$$

Note that the j^{th} coordinate of $\text{ALLOC}(\mathcal{D}, p)$ is the probability that a buyer with valuation $v \sim \mathcal{D}$ purchases item j under pricing p ; we will denote this probability by $\text{ALLOC}_j(\mathcal{D}, p)$.

Item Pricings and Sequential Item Pricings. For much of this work we will focus on *item pricing* mechanisms. For a single buyer, we represent an item pricing as a (potentially random) vector $p \in \mathbb{R}_+^m$ mapping individual items to prices, and extend it to sets additively by letting $p(S) := \sum_{j \in S} p_j$ for $S \subseteq [m]$. For a single buyer, we denote by ITEMPRICING the class of all randomized item pricing mechanisms, meaning that $p \in \mathbb{R}_+^m$ may be chosen at random.

In the multi-buyer setting we will further consider *Sequential Item Pricing* mechanisms. These mechanisms interact with buyers in a fixed order⁷. Let S_i denote the set of items left over after the mechanism has interacted with buyers 1 through $i - 1$. When buyer i arrives, the mechanism presents an item pricing over items S_i . The buyer instantiates his value and purchases his favorite bundle of items under the pricing, and then the mechanism moves on to the next buyer. The pricing presented to buyer i may depend on the instantiations of values of buyers 1 through $i - 1$ as well as the set S_i , but does not depend on the values of buyers $i \cdots n$.

Buy Many Mechanisms. In our results, we will also briefly refer to single-buyer *buy-many* mechanisms as defined by Chawla et al. [2022b]. Informally, a price menu p is buy-many if for every adaptive buyer strategy for sequentially purchasing a sequence of lotteries, there exists a single lottery in the menu that is cheaper and allocates more items to the buyer. We refer the reader to [Chawla et al., 2022b] for a formal definition.

B General Monotone Valuations

In this section, we focus on general monotone valuation functions. We propose a sequential item pricing to achieve an $O(\min\{n, m^2\})$ approximation to ex-ante item pricing revenue, and we show that there are no item pricing mechanisms that achieve a revenue of $o(\min\{n, \sqrt{m}\})$ of ex-ante revenue.

⁶For a groundset X , we use the notation Δ^X to denote the set of all probability distributions over X .

⁷Our results work for an arbitrary given order over buyers, although in some contexts it may be beneficial for the seller to choose a particular order.

B.1 Upper Bound

Theorem 18. *Let \mathcal{D} be any joint distribution for n buyers and m items over general monotone valuation functions. Then*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \leq \min\{n, 4m^2\} \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING})$$

Proof. Let (x^*, p^*) be the optimal solutions achieving $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$, with x_i^* denoting the ex ante allocation constraint for buyer i and p_i^* denoting the corresponding distribution over item pricings. For convenience, let

$$\text{OBJ} = \text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) := \sum_{i=1}^n \text{REV}(\mathcal{D}_i, p_i^*).$$

We will first show how to obtain the $O(n)$ upper bound. This is in fact trivial. We first sample a buyer $i \in [n]$ uniformly at random, and then present them with an item pricing that is instantiated from p_i^* , while ignoring all other buyers $i' \neq i$. This will recover a revenue of precisely $\text{REV}(\mathcal{D}_i, p_i^*)$ and then the proof is completed from the fact that we chose buyer i uniformly at random and thus we clearly obtain a n -approximation.

Up next, we show how to obtain a $O(m^2)$ upper bound. Recall that since we are restricted to item pricings,

$$\text{OBJ} = \mathbb{E}_{p^*} \left[\sum_{i=1}^n \sum_{j=1}^m p_{ij} \cdot \text{ALLOC}_j(\mathcal{D}_i, p_i) \right],$$

and thus, there exists some item $j^* \in [m]$ for which

$$\mathbb{E}_{p^*} \left[\sum_{i=1}^n p_{ij^*} \cdot \text{ALLOC}_{j^*}(\mathcal{D}_i, p_i) \right] \geq \frac{\text{OBJ}}{m}$$

For a fixed buyer i , consider the (random) uniform item pricing that sets the price of *each* item $j \in [m]$ to p_{ij^*}/m , given a price p_i that is instantiated from p_i^* . Observe that for any realization $v_i \sim \mathcal{D}_i$ where the ex ante buyer allocates to item j^* under pricing p_i , the bundle being allocated (and by monotonicity, also the set of all items $[m]$) has value at least p_{ij^*} . Therefore, since under price vector q_i the cost for buying all the items is precisely p_{ij^*} , buyer v_i will buy at least one item (assuming that all items are available), as it would rather buy the entire bundle rather than buy nothing.

In effect, the probability of buyer i buying at least one item under pricing q_i is $\lambda_i \geq \text{ALLOC}_{j^*}(\mathcal{D}_i, p_i)$ whenever all the items are available. Therefore, by rejecting to sell anything (via setting q_i to be infinity) with probability $1 - \frac{\text{ALLOC}_{j^*}(\mathcal{D}_i, p_i)}{2\lambda_i}$, we can ensure that the buyer i buys at least one item with probability $\frac{\text{ALLOC}_{j^*}(\mathcal{D}_i, p_i)}{2}$, netting a revenue of at least p_{ij^*}/m when doing so.

Finally, we go through each buyer i one-by-one, selling at price q_i with rejection as defined if all m items are available, and not selling anything if some item has already been sold. Notice that since we sample p_i from p_i^* , the probability that a buyer i doesn't purchase anything is precisely $x_{ij^*}^*/2$ under our pricing. To complete the argument, it remains to argue that whenever each buyer $i \in [n]$ arrives, there is a constant probability that all the items are available. As $\sum_{i=1}^n x_{ij^*}^* \leq 1$, by an OCRS argument similar to Lemma 6, all items are available with probability at least $1/2$ upon reaching any buyer; hence, the total revenue extracted from this sequential item pricing is at least

$$\mathbb{E}_{p \sim p^*} \left[\sum_{i=1}^n \frac{1}{2} \cdot \frac{p_{ij^*}}{m} \cdot \frac{\text{ALLOC}_{j^*}(\mathcal{D}_i, p_i)}{2} \right] \geq \frac{\text{OBJ}}{4m^2}.$$

The proof is completed by comparing n to $4m^2$ and picking the better sequential item pricing out of the two. \square

B.2 Lower Bound

We now present an instance that shows one cannot achieve an approximation of factor better than $\min\{n, \sqrt{m}\}$ of ex-ante revenue through sequential item pricing for general monotone valuations.

Theorem 19. *There exists a joint distribution \mathcal{D} for n buyers and m items over general monotone valuations for which*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \geq \Omega(\min\{n, \sqrt{m}\}) \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

As with the $\sqrt{\log m}$ lower bound for XOS valuations in Section 8, this lower bound on the gap between ex ante and sequential item pricing for the class of monotone valuations is more about the difficulty of finding an integral partition of the items that satisfies a typical realization of the buyers' valuation functions than it is about doing so in a sequential, item pricing setting. In particular, the family of instances which demonstrate Theorem 19 will show the same gap for EXPOSTITEMPRICING as well, which is an even stronger separation.

This gap is even more precipitous in the sense that it is purely combinatorial, and does not rely on the strategic limitations of mechanisms at all. In particular, this gap holds even for a “clairvoyant” mechanism which is given direct access to the realizations $v_i \sim \mathcal{D}_i$ and can extract the full surplus of the buyers' valuations of their allocations as revenue.

Before constructing our instance, we will introduce the key property which we require of it.

Definition 6. Let $m = \ell^2$ be the number of items. Let $P = (B_0, \dots, B_{\ell-1})$ be a partition of these m items to ℓ groups of ℓ items each, so that we have $|B_i| = \ell$ and $B_i \cap B_j = \emptyset$ for all $i \neq j$. Let \mathcal{P} be a collection of such partitionings. We call a collection \mathcal{P} a *good collection* if for every two partitionings $P_a = (A_0, \dots, A_{\ell-1}), P_b = (B_0, \dots, B_{\ell-1}) \in \mathcal{P}$, we have that $A_i \cap B_j \neq \emptyset$ for all $i, j \in [\ell]$.

In other words, a good collection is one that for any two partitions P_a and P_b and any two bundles from them, the intersection of the bundles is non-empty. We would like to assign to each buyer such partition, so that if a buyer purchases one bundle, other buyers would not be incentivized to buy any items.

Lemma 20. *For any prime number ℓ and $m = \ell^2$, there exists a good collection of size ℓ .*

Proof. Set the ℓ^2 items in a $\ell \times \ell$ matrix, such that each item is indexed by a tuple of form (i, j) where $0 \leq i, j \leq \ell - 1$. We define partition $P_i = (B_{i0}, \dots, B_{i(\ell-1)})$ as follows:

$$B_{ij} = \{(x, y) : y = (xi + j) \bmod \ell\},$$

where (x, y) here is the index of the corresponding items. Now first note that for each $i \in \{0, 1, \dots, \ell - 1\}$, each item (x, y) belongs to exactly one bundle B_{ij} , and each bundle consists of exactly ℓ items. So the defined P_i 's indeed create a collection of partitions.

Next, we need to prove that for any $i, i', j, j' \in \{0, 1, \dots, \ell - 1\}$ with $i \neq i'$, bundles B_{ij} and $B_{i'j'}$ intersect. Equivalently, we need an item (x, y) for which

$$y \equiv xi + j \equiv xi' + j' \pmod{\ell}$$

As \mathbb{Z}_ℓ is a field due to ℓ being prime, combined with $i' - i \neq 0$, there exists some $x \in \{0, 1, \dots, \ell - 1\}$ where $x(i' - i) \equiv j' - j \pmod{\ell}$. Setting $y = (xi + j) \bmod \ell$ concludes the proof. \square

Lemma 21. *For any positive integer m , there exists a good collection of size at least $\ell = \frac{\lfloor \sqrt{m} \rfloor}{2}$ that uses ℓ^2 items.*

Proof. Due to Bertrand's postulate, we know that for any $n > 3$, there exists at least one prime number p such that $2n - 2 > p > n$. Let $n = \frac{\lfloor \sqrt{m} \rfloor + 2}{2}$. Therefore, there exists a prime number p such that $\lfloor \sqrt{m} \rfloor > p > \frac{\lfloor \sqrt{m} \rfloor + 2}{2}$. Now using Lemma 20, we know that we can build a collection of size p since p is a prime, and we know that $p > \frac{\lfloor \sqrt{m} \rfloor + 2}{2} > \frac{\lfloor \sqrt{m} \rfloor}{2}$. This concludes the proof. \square

Proof of Theorem 19. Let ℓ be the closest prime number no more than \sqrt{m} and let $N = \min(n, \ell)$. We prove the theorem in 3 parts: First, we show how to build the instance's valuation function. Then, we show that no ex-post mechanism can gain revenue of more than 1 under this valuation. Finally, we show that there exists an ex-ante pricing under which the revenue is at least $N(1 - \epsilon)$. This concludes the proof of theorem.

Instance. We index both the items and buyers from 0. Let $\mathcal{P} = \{P_0, \dots, P_{\ell-1}\}$ be the *good collection* on the first ℓ^2 items from Lemma 20. For $P_i = (B_{i1}, \dots, B_{i(\ell-1)s})$, let the valuation of buyer i on the m items be deterministic and as follows:

$$v_i(S) = \mathbb{1}\left[B_{ib} \subseteq S \text{ for some } b \in \{0, 1, \dots, \ell-1\}\right] \text{ for } 0 \leq i \leq \ell-1, \quad v_i(S) = 0 \text{ for } i \geq N$$

In other words, we have set zero value for all but the first ℓ^2 items, and any buyer with index $i \geq N$ has zero valuation on any set of items. Moreover, buyer $i \in \{0, 1, \dots, N-1\}$ values any set containing one of their bundles by 1, and otherwise 0. This is a monotone valuation, but not sub-additive.⁸

Upper bound on revenue of a mechanism. Using the property of the *good collection*, if an ex-post mechanism sells a set S containing one of their bundles to a buyer, all the other buyers will value the remaining sets of items by zero. This is because the set S contains at least one item from any bundle of any of the other buyers, and a buyer values a set by zero if their set does not include one of their whole bundles. This means that any mechanism—sequential item pricing mechanisms included—can sell to at most one buyer. Since the maximum social welfare of a buyer is 1, the maximum revenue achievable will also be at most 1, and can be gained by selling one bundle to a buyer.

Lower bound on ex ante revenue. Up next, consider the following mechanism. For each of the buyers $i \in \{0, 1, \dots, N-1\}$, pick one of their bundles uniformly at random. This will happen with probability $\frac{1}{\ell}$ for each bundle. Price all of the items in this bundle at $\frac{1-\epsilon}{\ell}$ for some $\epsilon > 0$, and all other items at infinity. In this case, the buyer purchases the randomly selected bundle, and gains a utility of $1 - \ell \cdot \frac{1-\epsilon}{\ell} = \epsilon > 0$. The revenue from this draw and this buyer will be $1 - \epsilon$. Now note that each item is sold to each buyer with probability $1/\ell$ (i.e. the probability that its corresponding bundle is selected) and so the total allocation of item j is $N/\ell \leq 1$. Therefore, the ex-ante constraint is satisfied, and this mechanism gets a total revenue of $N \cdot (1 - \epsilon)$.

We can conclude that for this instance, no ex-post mechanism can get a revenue of better than a factor of N of ex-ante revenue. Now using Lemma 21, we know that $\ell = \Omega(\sqrt{m})$. This concludes the bound in the theorem statement. \square

C Gross Substitutes Valuations

The analysis presented in the main part of this work establishes the existence of a sequential item pricing mechanism that achieves a $4e/(e-1)$ -approximation to $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$ whenever \mathcal{D} is a distribution over gross substitutes (henceforth, GS) valuations. This was achieved through the design of a 1-RRS for GS valuations (Proposition 8), which in turn implies a $e/(e-1)$ -OCRS (Lemma 7) and finally a $4e/(e-1)$ -approximation of $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$ (Lemma 6).

However, the previous work of Chawla et al. [2023] established a 2-approximation for *unit demand* valuations (which are also GS), and furthermore this constant is easily shown to be tight for unit-demand valuations even for $m = 1$ items via standard prophet inequality bounds. In this section, we will show that one can obtain this factor of 2 even for GS valuations, proving Theorem 5 which we restate for convenience.

Theorem 5. *Let \mathcal{D} be any joint distribution for n buyers and m items over gross substitutes valuation functions. Then we have,*

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) \leq 2 \cdot \text{REV}(\mathcal{D}, \text{SEQITEMPRICING}).$$

⁸Let B_{ib} be a bundle of buyer i . Then for any $S \subsetneq B_{ib}$, $|S| > 0$ we have that $v_i(S) + v_i(B_{ib} \setminus S) = 0 < v_i(B_{ib}) = 1$.

In Section C.1 we argue that the analysis of CONVEXHULLSAMPLER can be tightened under GS valuations to show that we can actually turn the 1-RRS for GS valuations into a 1-OCRS, removing the loss of $e/(e-1)$. This already implies a 4-approximation for $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$, through Lemma 6. Unfortunately, we cannot further reduce the approximation constant, as it is inherited from the generality with which we have stated our definitions, in order to capture harder families of valuations.

In Section C.2 we replicate the analysis of Chawla et al. [2023], in order to obtain a 2-approximation for $\text{EAREV}(\mathcal{D}, \text{ITEMPRICING})$ under any GS valuation \mathcal{D} and prove Theorem 5. While the proof in [Chawla et al., 2023] was only shown for unit demand valuations, we observe that the same arguments also hold for GS valuations. We comment that while the idea in this proof is similar to the ones used in this paper, it heavily relies on the GS property and thus, cannot be directly applied to the more general families of valuations that this work addresses.

C.1 Analysis of ConvexHullSampler for GS valuations

In this section, we show that the $(1 - 1/e)$ -loss in revenue from the analysis of CONVEXHULLSAMPLER in Section 4 is actually not necessary under GS valuations. First, recall that when we defined our 1-RRS for GS valuations, we simply set $q := p$ as the GS condition handles the revenue constraint (see the paragraph above Proposition 8 for the full argument). Let the $\{y^T\}$ and w vectors be defined exactly as in the application of CONVEXHULLSAMPLER in Section 4. Then, observe that by the definition of the $\{y^T\}$ vectors we have that for any set $Q \subseteq [k]$ and any $j \in Q$:

$$y_j^Q = \alpha \cdot q_j^Q \text{ALLOC}_j(\mathcal{D}_{|Q}, q^Q) = p_j \cdot \text{ALLOC}_j(\mathcal{D}_{|Q}, p) \geq p_j \cdot \text{ALLOC}_j(\mathcal{D}, p) = w_j \quad (8)$$

where the inequality crucially uses the fact that \mathcal{D} is GS, and thus the allocation of an item under any restriction (that includes it) will never decrease. This inequality is all we need in order to provide a better analysis of CONVEXHULLSAMPLER that doesn't suffer the $(1 - 1/e)$ -loss.

Specifically, we proceed to analyze CONVEXHULLSAMPLER exactly as in the proof of Lemma 10, with the only difference that instead of proving that

$$\sum_{T \subseteq S} \lambda_T \cdot |y^T| \geq \left(1 - \frac{1}{e}\right) \cdot |w|,$$

we will actually show that

$$\sum_{T \subseteq S} \lambda_T \cdot |y^T| \geq |w|,$$

from which it becomes clear that the $(1 - 1/e)$ -factor is no longer lost. To prove this, let s be the number of iterations of CONVEXHULLSAMPLER and recall that either $\sigma = 1$ or $Q_s = \emptyset$ from our termination condition. We will show that (8) implies that $Q_s = \emptyset$ necessarily. Indeed, assume that this is not the case, and let $j \in Q_s$. Since $Q_s \neq \emptyset$, it must be the case that $\sigma = 1$ and thus $\sum_{i=0}^{s-1} \lambda_{Q_i} = 1$. Since $j \in Q_s$, we have that

$$0 < \tilde{w}_{s,j} := w_j - \sum_{i=0}^{s-1} \lambda_{Q_i} \cdot y_j^{Q_i}$$

which is clearly a contradiction from (8), as the sum is simply a linear combination of terms that are at least w_j . Thus, we have concluded that $Q_s = \emptyset$ always, which means that $x_s = 0$ and thus $w = \sum_{T \subseteq S} \lambda_T \cdot y^T$, which is precisely what we wanted to show.

C.2 A 2-approximation for GS valuations

In this section, we replicate the analysis of Chawla et al. [2023] to prove Theorem 5. The proof is centered around the following lemma, which essentially combines the *Revenue Recovery Scheme* and the *Convex Hull Sampling* steps of our approach.

Lemma 22. Fix any GS valuation \mathcal{D} , any distribution over subsets of items \mathcal{S} , any deterministic item pricing p , and let $w := \mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{ALLOC}(\mathcal{D}_{|\mathcal{S}}, p)]$. Then, for any vector $y \preceq w$ there exists a randomized item pricing $q = q(\mathcal{D}, \mathcal{S}, p, y)$ such that:

- (a) $\mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{ALLOC}(\mathcal{D}_{|\mathcal{S}}, q)] = y$, and,
- (b) $\mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{REV}(\mathcal{D}_{|\mathcal{S}}, q)] = \sum_{j=1}^m p_j \cdot y_j$.

Proof. For each set of items $T \subseteq [m]$, we define the vector $x^T := \mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{ALLOC}(\mathcal{D}_{|S \cap T}, p)]$ and observe that $x_j^T = 0$ for any $j \notin T$ and $x_j^T \geq w$ for any $j \in T$, due to the GS property of \mathcal{D} . Thus, since $y \preceq w$, we can immediately deduce that vector y lies in the convex hull of the x^T vectors, or equivalently that there exist non-negative coefficients $\{\lambda_T\}_{T \subseteq [m]}$ such that $\sum_{T \subseteq [m]} \lambda_T = 1$ and $y = \sum_{T \subseteq [m]} \lambda_T \cdot x^T$.

The proposed item pricing q samples a set $T \subseteq [m]$ according to the probabilities λ_T , and then presents the buyer with pricing q^T , defined as $q_j^T = p_j$ if $j \in T$ and $q_j^T = \infty$ otherwise. We will now argue that this pricing satisfies the two conditions of the claim.

For the expected allocation, we have

$$\mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{ALLOC}(\mathcal{D}_{|\mathcal{S}}, q)] = \sum_{T \subseteq [m]} \lambda_T \cdot \mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{ALLOC}(\mathcal{D}_{|\mathcal{S}}, q^T)] = \sum_{T \subseteq [m]} \lambda_T \cdot \mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{ALLOC}(\mathcal{D}_{|S \cap T}, p)] = y,$$

with the first equality following from q 's definition, the second equality follows from the fact that q^T is equivalent to p under valuations restricted on T , and the last equality follows from the definition of x^T and the fact that $y = \sum_{T \subseteq [m]} \lambda_T \cdot x^T$.

Likewise, for the expected revenue, with similar arguments we have

$$\mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{REV}(\mathcal{D}_{|\mathcal{S}}, q)] = \sum_{T \subseteq [m]} \lambda_T \sum_{j=1}^m p_j \mathbb{E}_{\mathcal{S} \sim \mathcal{S}}[\text{ALLOC}_j(\mathcal{D}_{|S \cap T}, p)] = \sum_{j=1}^m p_j \cdot \left(\sum_{T \subseteq [m]} \lambda_T x_j^T \right) = \sum_{j=1}^m p_j \cdot y_j.$$

□

We are now ready to prove Theorem 5. Consider any GS valuations $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ and recall that

$$\text{EAREV}(\mathcal{D}, \text{ITEMPRICING}) = \sum_{i=1}^n \text{EAREV}_{x_i^*}(\mathcal{D}_i, \text{ITEMPRICING})$$

for some ex-ante constraint $\{x_i^*\}_{i=1}^n$ with $\sum_{i=1}^n x_{ij}^* \leq 1$ for all items $j \in [m]$. Consider the sequential item pricing mechanism that offers each buyer $i \in [n]$ the pricing $q_i = q(\mathcal{D}_i, \mathcal{S}_i, p_i, y_i)$ of Lemma 22 over only the available items, for the following instantiations: \mathcal{D}_i is the distribution of valuations for buyer i ; \mathcal{S}_i is the distribution of available items when i arrives (that depends on the randomness of pricings and allocations over previous buyers); p_i is an item pricing sampled from the distribution that achieves $\text{EAREV}_{x_i^*}(\mathcal{D}_i, \text{ITEMPRICING})$; and $y_i = (1/2) \cdot \text{ALLOC}(\mathcal{D}_i, p_i)$.

Define $w_i := \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}_i}[\text{ALLOC}(\mathcal{D}_{|\mathcal{S}_i}, p_i)]$. As long as we can prove that $y_i \preceq w_i$ for all $i \in [n]$, then the guarantees of Lemma 22 will hold, and from condition (b) we will immediately obtain that the revenue the above sequential item pricing collects from each buyer i (call it REV_i) will satisfy

$$\text{REV}_i = \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}_i, p_i}[\text{REV}(\mathcal{D}_{|\mathcal{S}_i}, q_i)] = \frac{1}{2} \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}_i, p_i} \left[\sum_{j=1}^m p_j \cdot \text{ALLOC}_j(\mathcal{D}_i, p_i) \right] = \frac{1}{2} \cdot \text{EAREV}_{x_i^*}(\mathcal{D}_i, \text{ITEMPRICING}),$$

and thus the proof of Theorem 5 will be completed.

In order to argue that $y_i \preceq w_i$, observe that for any item $j \in [m]$ we have

$$w_{ij} = \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}_i}[\text{ALLOC}_j(\mathcal{D}_{|\mathcal{S}_i}, p_i)] = \Pr[j \in \mathcal{S}_i] \cdot \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}_i}[\text{ALLOC}_j(\mathcal{D}_{|\mathcal{S}_i}, p_i) | j \in \mathcal{S}_i] \geq \Pr[j \in \mathcal{S}_i] \cdot \text{ALLOC}_j(\mathcal{D}_i, p_i),$$

where the inequality follows from the GS property, under which the allocation of any item under any restriction that includes it can only increase. Thus, a sufficient condition to show that $y_i \preceq w_i$ is to prove that $\Pr[j \in S_i] \geq \frac{1}{2}$ for all buyers $i \in [n]$. This will certainly be true for the first buyer, as $S_1 = [m]$. Then, from condition (a) of Lemma 22 we obtain that the expected allocation of items to the first buyer is exactly y_1 . Inductively, we proceed to show that each buyer $i' < i$ gets expected allocation $y_{i'}$. Then, we get that

$$\Pr[j \in S_i] = 1 - \sum_{i' < i} y_{i'j} \geq 1 - \sum_{i'=1}^n y_{i'j} = 1 - \frac{1}{2} \sum_{i'=1}^n x_{i'j}^* \geq \frac{1}{2},$$

with the last inequality following from the ex ante constraint. Thus, we have shown via induction that $\Pr[j \in S_i] \geq 1/2$ for all $j \in [m]$ and $i \in [n]$, completing the proof.

D A Lower Bound for RRSes on XOS Valuations

In Lemma 9, we proved the existence of an $O(\log m + \log \Gamma)$ -RRS, where $\Gamma := p_{\max}/p_{\min}$ is the aspect ratio of the given price. We now show that this dependency on Γ is necessary: without it, we would have to suffer a factor of $\Omega(\sqrt{m})$ in our approximation. This is captured through the following.

Theorem 23. *For all m , there exists a distribution \mathcal{D} over XOS valuations, a deterministic item pricing p , and a deterministic subset of items $S \subseteq [m]$, such that for all item pricings q ,*

$$\text{REV}(\mathcal{D}_{|S}, q) \leq \frac{1}{\Omega(\sqrt{m})} \cdot \sum_{j \in S} p_j \cdot \text{ALLOC}_j(\mathcal{D}, p).$$

This rules out the existence of any $o(\sqrt{m})$ -RRS for XOS valuations.

Before proceeding, we note that this lemma is not a counterargument to Informal Theorem 1, as the ex ante allocation constraints for the counterexample are indeed exponential in m . Our counterexample is as follows.

Available items and pricings. We define the subset of items $S := [m-1]$ (i.e. only the last item is not available), and the pricing vector $p := (\beta, \beta^2, \dots, \beta^{m-1}, 0)$, where $\beta := \sqrt{m-1}$.

Valuation distribution. We define the (deterministic) XOS valuation function $v_{i,R}$ that is parametrized by an index $i \in [m-1]$ and a set $R \subseteq [i-1]$. Such a valuation function consists of two additive components $v_{i,R}^1$ and $v_{i,R}^2$, i.e.

$$v_{i,R}(T) := \max \left\{ \sum_{j \in T} v_{i,R}^1(j), \sum_{j \in T} v_{i,R}^2(j) \right\}$$

for any $T \subseteq [m]$, where the additive components are defined as

$$v_{i,R}^1(j) = \begin{cases} \beta^i & \text{if } j = i, \\ \epsilon + \sum_{k \in R} (\beta^i - \beta^k) & \text{if } j = m, \\ 0 & \text{otherwise.} \end{cases} \quad \text{and} \quad v_{i,R}^2(j) = \begin{cases} \beta^i & \text{if } j \in R, \\ 0 & \text{otherwise.} \end{cases}$$

for some $\epsilon > 0$. The buyer valuation distribution \mathcal{D} is then defined via the following sampling procedure:

1. An index $i \in [m-1]$ is sampled proportionally to β^{-i} , i.e. with probability $\frac{\beta^{-i}}{\sigma}$ for $\sigma := \sum_{i=1}^{m-1} \beta^{-i}$.
2. A set $R \subseteq [i-1]$ is sampled by including each element $j \in [i-1]$ with probability $(m-1)^{-1/2}$ independently.

3. Finally, valuation $v_{i,R}$ is realized.

We now begin the proof by examining the right-hand side of the statement of Theorem 23. Observe that for any $i \in [m-1]$ and $R \subseteq [i-1]$, the utility-maximizing set under the valuation function $v_{i,R}$ and our chosen pricing p is the set $\{i, m\}$; note that this is the utility-maximizing set for the first additive component with utility of $\epsilon + \sum_{k \in R} (\beta^i - \beta^k)$, while the maximum utility from the second additive component is $\sum_{k \in R} \beta^i - \beta^k$. Therefore, we have $\text{ALLOC}_j(\mathcal{D}, p) = \beta^{-j}/\sigma$ for all $j \in S$ (as whenever we draw j in the sampling process of \mathcal{D} , we know automatically that the utility-maximizing set is $\{j, m\}$). Hence,

$$\sum_{j \in S} p_j \cdot \text{ALLOC}_j(\mathcal{D}, p) = \sum_{j=1}^{m-1} \beta^j \cdot \frac{\beta^{-j}}{\sigma} = \sigma^{-1} \cdot (m-1).$$

We also comment that the allocation probabilities are exponentially small in m , and thus they do not satisfy the assumption of Informal Theorem 1.

To complete the proof, it suffices to show that under any pricing q , we have

$$\text{REV}(\mathcal{D}_{|S}, q) \leq 4\sigma^{-1}\sqrt{m-1}. \quad (9)$$

Fix any pricing q . We define a labelling vector $\ell \in \mathbb{Z}^{m-1}$, where $\ell_j = \lceil \log_\beta q_j \rceil$ for all $j \in [m-1]$. We now consider the maximum revenue we can extract from $v_{i,R|S}$ for any i and R . Since $v_{i,R}$ is defined as the maximum over two additive valuations, we consider the three following ways via which we can extract revenue from it.

Revenue is extracted from $v_{i,R|S}^1$ and $\ell_i \neq i$. Since revenue is extracted from $v_{i,R|S}^1$ and $m \notin S$, the only item that can be purchased is i , valued at β^i . Furthermore, since $\ell_i \neq i$, we can conclude that $q_i \leq \beta^{i-1}$ and therefore, the total revenue extracted from this case is at most

$$\mathbb{E}_{i,R}[\beta^{i-1}] = \sum_{i=1}^{m-1} \beta^{i-1} \cdot \frac{\beta^{-i}}{\sigma} = \sigma^{-1} \cdot \frac{m-1}{\beta} = \sigma^{-1}\sqrt{m-1}.$$

Revenue is extracted from $v_{i,R|S}^1$ and $\ell_i = i$. In this case, the revenue extracted is at most β^i . A critical observation here is that we need $\ell_j \geq i$ for all $j \in R$ in order for this case to happen; note that the utility of S under $v_{i,R|S}^1$ is $\beta^i - q_i < \beta^i - \beta^{i-1}$, since $\ell_i = i$, while if any item $j \in R$ has $\ell_j \leq i-1$, then its utility under $v_{i,R|S}^2$ is $\beta^i - q_j \geq \beta^i - \beta^{i-1}$, and thus the buyer would deviate away from buying item i and towards buying this j under $v_{i,R|S}^2$.

We now define the set $U = \{i \in S : \ell_i = i\}$, and note that revenue in this case can only be extracted from items in U . Furthermore, consider any item $i \in U$ and any R , we can only extract revenue from item i under $v_{i,R|S}^1$ if $R \cap U = \emptyset$; otherwise, some item $j \in R$ has $\ell_j = j < i$, and by the previous observation we cannot extract revenue from item i under $v_{i,R|S}^1$. Therefore, if we enumerate $U = \{u_1, u_2, \dots, u_t\}$ where $u_1 < u_2 < \dots < u_t$, the total revenue extracted in this case is at most

$$\begin{aligned} \mathbb{E}_{i,R}[\beta^i \cdot \mathbb{1}(\text{revenue from } v_{i,R|S}^1 \text{ and } \ell_i = i)] &\leq \mathbb{E}_{i,R}[\beta^i \cdot \mathbb{1}[i \in U \wedge R \cap U = \emptyset]] \\ &= \sum_{j=1}^t \beta^j \Pr[i = u_j] \Pr[R \cap U = \emptyset \mid i = u_j] \\ &= \sum_{j=1}^t \beta^{u_j} \cdot \frac{\beta^{-u_j}}{\sigma} \cdot \left(1 - \frac{1}{\sqrt{m-1}}\right)^{j-1} \\ &\leq \sigma^{-1} \sum_{j=0}^{\infty} \left(1 - \frac{1}{\sqrt{m-1}}\right)^j = \sigma^{-1}\sqrt{m-1} \end{aligned}$$

where the second line uses the fact that each item u_1, u_2, \dots, u_{j-1} does not appear in R with probability $1 - \frac{1}{\sqrt{m-1}}$ independently.

Revenue is extracted from $v_{i,R|S}^2$. We first fix any item j , and for convenience let $l = \ell_j$. Observe that item j can only be sold in this case if $l \leq i$ (otherwise when $l > i$, we have $q_j > \beta^{l-1} \geq \beta^i = v_{i,R|S}^2(j)$); furthermore, we obviously have $q_j \leq \beta^l$. Therefore, the total revenue extracted from item j in this case is at most

$$\begin{aligned} \mathbb{E}_{i,R} [\beta^l \cdot \mathbb{1}(\text{revenue from } v_{i,R|S}^2 \text{ and } j \text{ is included})] &\leq \mathbb{E}_{i,R} [\beta^l \cdot \mathbb{1}[j \in R \wedge i \geq l]] \\ &= \beta^l \sum_{k=1}^{m-1} \Pr[i = k] \Pr[j \in R \mid i = k] \\ &\leq \frac{\beta^l}{\sqrt{m-1}} \sum_{k=l}^{m-1} \frac{\beta^{-k}}{\sigma} \\ &\leq \frac{\beta^l}{\sqrt{m-1}} \cdot \frac{2\beta^{-l}}{\sigma} = \frac{2\sigma^{-1}}{\sqrt{m-1}} \end{aligned}$$

where the second line uses that $\Pr[j \in R \mid i = k]$ is $\frac{1}{\sqrt{m-1}}$ if $i > j$, or 0 otherwise; the second line assumes that $\beta = \sqrt{m-1} \geq 2$, which means we can bound $\sum_{k=l}^{m-1} \beta^{-k} \leq 2\beta^{-l}$.

Thus, the total revenue extracted from this case is at most

$$\mathbb{E}_{i,R} \left[\sum_{j=1}^m \beta^{\ell_j} \cdot \mathbb{1}(\text{revenue from } v_{i,R|S}^2 \text{ and } j \text{ is included}) \right] \leq 2\sigma^{-1}\sqrt{m-1}.$$

Combining the three cases gives us an upper bound of $4\sigma^{-1}\sqrt{m-1}$ on $\text{REV}(\mathcal{D}_{|S}, q)$ for any q .

E Deferred Proofs from Section 8

We start by reproducing the description of our instance for convenience.

The buyers. Let n buyers $v_i \sim \mathcal{D}_i$ be (independently) identically distributed, with valuations sampled according to the following process:

1. A set of items $A \subseteq [m]$ with $|A| = k$ is sampled uniformly at random.
2. An integer $h \in \{1, 2, \dots, \frac{1}{2} \log k\}$ is drawn uniformly at random. We denote $\ell := 2^h$.
3. The valuation of the buyer is then realized as an XOS function, which is given by the maximum over the following additive valuations:
 - (a) For the set A , we define an additive valuation v_A such that $v_A(j) = 1 + t$ for items $j \in A$ and $v_A(j) = 0$ for items $j \notin A$.
 - (b) For *every* set B of cardinality $|B| = t \cdot \ell$, we define an additive valuation v_B such that $v_B(j) = 1 + \frac{k}{\ell}$ for items $j \in B$ and $v_B(i) = 0$ for items $j \notin B$.

We will take $n = k = \sqrt{m}$ and $t = \sqrt{\log k}$. Observe that for our selected parameters $t \cdot \ell \leq \sqrt{m \cdot \log m} \leq m$, and all parameters are integer-valued, so this is a well-defined instance. Let A_i be the A -set for buyer i .

In order to prove the instance-specific claims in Section 8, we will require more notation. We use $\text{REV}(v_i, p)$ to denote the revenue that we collect from the i -th buyer if their valuation is realized to v_i and the pricing is some fixed vector $p \in \mathbb{R}_{\geq 0}^m$. Note that this is a deterministic quantity, as all sources of randomness are fixed.

For our instance we will define the events $\mathcal{E}^A(v_i, p)$ and $\mathcal{E}^B(v_i, p)$ that subset of items $[m]$ that buyer i with valuation v_i purchases is of type A or B respectively; more formally, $\mathcal{E}^A(v_i, p)$ holds if for prices p the buyers' utility for their preferred bundle $T \subseteq [m]$ is given by

$$u_i(T) = v_i(T) - p(T) = \sum_{j \in T} (v_A(j) - p_j),$$

and otherwise $\mathcal{E}^B(v_i, p)$ holds if their valuation is given by some v_B .

Lemma 14. *For any v_i, p such that $\mathcal{E}^A(v_i, p)$ holds, buyer i purchases at least $\frac{kt}{t+1}$ items.*

Proof of Lemma 14. We consider a fixed buyer $v = v_i$ (and thus also the set A and the integer ℓ in the valuation description) and a fixed pricing p . Since $\mathcal{E}^A(v, p)$ holds, we know that the buyer will purchase a set of items such that their utility is maximized through the linear valuation v_A ; in particular, including a subset $A_T \subseteq A$ given by

$$A_T := \arg \max_{T \subseteq A} (v_A(T) - p(T)) = \{j \in A : p_j \leq t + 1\}.$$

We will show that $|A_T| \geq \frac{kt}{t+1}$. To that end, let $B_{AT} \subseteq A_T$ be a maximal B -utility subset of A_T . Observe that since $t \leq \frac{k}{\ell}$ for our chosen parameters, every item worth buying under an A -valuation v_A will also be worth buying under some B -valuation v_B . Thus, B_{AT} is any maximal-cardinality subset of A_T of size at most $t\ell$.

Since $\mathcal{E}^A(v_i, p)$ holds we know that i prefers A_T to B_{AT} , meaning A_T has higher utility than B_{AT} , so

$$\sum_{j \in A_T} (t + 1 - p_j) \geq \sum_{j \in B_{AT}} \left(1 + \frac{k}{\ell} - p_j\right). \quad (10)$$

We first argue that $|A_T| > t\ell$. Otherwise $B_{AT} = A_T$, and thus (10) can only be satisfied if $t \geq \frac{k}{\ell}$ which is not the case as $t\ell \leq \sqrt{\log k} \cdot \sqrt{k} < k$. Therefore $|A_T| \geq t\ell$. In this case, $|B_{AT}| = t\ell$ and (10) can be restated as

$$\sum_{j \in A_T \setminus B_{AT}} (t + 1 - p_j) \geq \sum_{j \in B_{AT}} \left(\frac{k}{\ell} - t\right) = kt - t^2\ell, \quad (11)$$

and since $p_j \geq 0$ for all j this implies that

$$(t + 1) \cdot (|A_T| - t\ell) \geq kt - t^2\ell.$$

This immediately implies that $|A_T| \geq \frac{kt}{t+1}$, as claimed. \square

Lemma 15. *For some $c_1 \in \mathbb{R}_{\geq 0}$ and sufficiently large k , for all p , $\mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^A(v_i, p)] \leq c_1 \cdot k$.*

Proof of Lemma 15. We will prove this claim for every item pricing p and every buyer valuation $v = v_i \sim \mathcal{D}_i$, provided that the event $\mathcal{E}^A(v_i, p)$ holds. The claim will then follow from taking the conditional expectation over v_i .

We use the same setup as in the proof of Lemma 14 to define the sets A_T and B_{AT} . Since $\mathcal{E}^A(v_i, p)$ holds by assumption, we know that $\text{REV}(v_i, p) = p(A_T)$ because the buyer maximizes their utility via v_A and buys this set. Therefore we will bound $p(A_T)$.

Just as before, (10) implies that $|A_T| \geq t\ell$ and that

$$\sum_{j \in A_T \setminus B_{AT}} (t + 1 - p_j) \geq kt - t^2\ell \quad (12)$$

holds for *any* subset B_{AT} of A_T with cardinality $t\ell$; in particular it holds for B_{AT} the $t\ell$ items in A_T of minimum price. For this choice of B_{AT} the average item price in $A_T \setminus B_{AT}$ exceeds the average item price in A_T , i.e.

$$\frac{p(A_T \setminus B_{AT})}{|A_T \setminus B_{AT}|} \geq \frac{p(A_T)}{|A_T|}.$$

Plugging this into (12) and observing that $|A_T \setminus B_{AT}| = |A_T| - t\ell$ yields

$$kt - t^2\ell \leq \sum_{j \in A_T \setminus B_{AT}} (1 + t - p_j) = (1 + t)(|A_T| - t\ell) - p(A_T \setminus B_{AT}) \leq (|A_T| - t\ell) \cdot \left(1 + t - \frac{p(A_T)}{|A_T|}\right).$$

Solving for $p(A_T)$ we obtain

$$p(A_T) \leq |A_T| \cdot \frac{|A_T|(t+1) - t(k+\ell)}{|A_T| - t\ell},$$

and using that $|A_T| \leq |A| = k$ (by definition) and that $|A_T| \geq \frac{kt}{t+1}$ by Lemma 14 since $\mathcal{E}_i^A(v, p, S)$ holds, we finally obtain that

$$p(A_T) \leq k \cdot \frac{k(t+1) - t(k+\ell)}{\frac{kt}{t+1} - t\ell}.$$

The proof is completed by showing

$$\frac{k(t+1) - t(k+\ell)}{\frac{kt}{t+1} - t\ell} \leq c_1$$

for some constant c_1 and all sufficiently large k , which easily follows from observing that $k \geq 2t\ell$ for our choice of parameters $t = \sqrt{\log k}$ and $\ell \in [2, \sqrt{k}]$. \square

Lemma 16. *For some $c_2 \in \mathbb{R}_{\geq 0}$ and sufficiently large k , for all p , $\mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p)] \leq c_2 \cdot t \cdot \frac{k}{\log k}$.*

Proof. Consider a fixed item pricing $p \in \mathbb{R}_{\geq 0}^m$. We can relabel the items so that their prices are non-decreasing, i.e. we assume without loss of generality that $p_1 \leq p_2 \leq \dots \leq p_m$.

Since we are conditioning on event $\mathcal{E}_i^B(v, p, S)$, we know that the buyer always buys a set of items whose utility is maximized through a B -valuation, thus they buy a set of at most $t\ell$ items where ℓ is a buyer-dependent parameter that is drawn uniformly at random from $L = \{2, 4, 8, \dots, \sqrt{k}\}$. Furthermore, the buyer will clearly prioritize the cheaper items, as all items have the same (additive) valuation of $1 + k/\ell$ under B and cheaper items contribute more to the buyer's utility. Thus, for a given ℓ , the buyer will purchase items T which are a prefix of the items in ascending-price-order:

$$T = \{1, 2, 3, \dots, x\} \subseteq [m],$$

where x is the maximum index for which $x \leq t\ell$ and $p_x \leq 1 + \frac{k}{\ell}$. We will bound the conditional expected revenue, which is precisely the conditional expectation of $p(T)$, by observing that

$$\begin{aligned} \mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p)] &= \sum_{\ell' \in L} \mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p) \wedge \ell(v_i) = \ell'] \cdot \Pr_{v_i} [\ell(v_i) = \ell'] \\ &= \frac{1}{|L|} \cdot \sum_{\ell' \in L} \mathbb{E}_{v_i} [p(T) \mid \mathcal{E}^B(v_i, p) \wedge \ell(v_i) = \ell']. \end{aligned} \quad (13)$$

Here we use $\ell(v_i)$ to denote the value of ℓ that $v_i \sim \mathcal{D}_i$ realizes, and we used that $\Pr_{v_i} [\ell(v_i) = \ell'] = \frac{1}{|L|} = \frac{2}{\log k}$ for all $\ell' \in L$. We will bound this sum (13) by $c_2 \cdot \frac{tk}{\log k}$ for some constant c_2 and sufficiently large k . We will do this by a combination of case analysis and grouping the possible values of ℓ as a function of the non-uniform prices. Let

$$\ell^* := \max \left(\left\{ \ell \in L : p_{t\ell} \leq \frac{k}{\ell} + 1 \right\} \right).$$

This set is nonempty and the threshold parameter value ℓ^* exists provided that $p_{2t} \leq \frac{k}{2} + 1$. This is because the thresholds $\frac{k}{\ell} + 1$ are decreasing in ℓ , while the prices $p_{t\ell}$ are increasing in ℓ . Therefore we have two cases to consider.

Case 1. First, we consider the case that $p_{2t} > \frac{k}{2} + 1$. In this case, since $t + 1 \leq \frac{k}{2} + 1$ for $t = \sqrt{\log k}$ and also $\frac{k}{\ell} + 1 \leq \frac{k}{2} + 1$ for $\ell \geq 2$, we know the buyer will not buy any item j for $j \geq 2t$; since p_j is too high it would necessarily decrease their utility. Furthermore, the maximum revenue attainable from each of the $2t$ remaining items is at most equal to their value to the buyer, which is $(\frac{k}{\ell} + 1)$ for $\ell \in L$ uniformly at random. Thus, from (13) we obtain

$$\mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p)] \leq 2t \sum_{\ell \in L} \frac{1}{|L|} \cdot \left(1 + \frac{k}{\ell}\right) = \frac{1}{|L|} \cdot (2t|L| + 2tk) \leq \frac{6tk}{\log k},$$

where the last statement follows from $|L| = \frac{1}{2} \log k$ and the fact that $L = \{2, 4, 8, \dots, \sqrt{k}\}$ follows a geometric progression and so the ℓ^{-1} terms sum to at most 1.

Case 2. We now address the case where ℓ^* is well-defined. We consider a buyer $v = v_i$ and again use $\ell(v)$ to denote the ℓ value that their valuation realized. We analyze their expected revenue based on whether $\ell(v) < \ell^*$, $\ell(v) = \ell^*$, or $\ell(v) > \ell^*$.

If $\ell(v) = \ell^*$, then we know that the buyer will purchase at most $t\ell^*$ items, and also that each item contributes at most the buyer's per-item value to the revenue, which is $\frac{k}{\ell^*} + 1$. Thus we have

$$\mathbb{E}_v [\text{REV}(v, p) \mid \mathcal{E}^B(v, p) \wedge \ell(v) = \ell^*] \leq t\ell^* \left(1 + \frac{k}{\ell^*}\right) = t(k + \ell^*). \quad (14)$$

If $\ell(v) < \ell^*$, then each item sold can contribute at most $(\frac{k}{\ell^*} + 1)$ to the revenue, and for each ℓ at most $t\ell$ items can be sold. Thus, we have

$$\sum_{\ell \in L: \ell < \ell^*} \mathbb{E}_v [\text{REV}(v, p) \mid \mathcal{E}^B(v, p) \wedge \ell(v) = \ell] \leq \sum_{\ell \in L: \ell < \ell^*} t\ell \cdot \left(1 + \frac{k}{\ell^*}\right) \leq t(k + \ell^*), \quad (15)$$

where the last inequality follows from $\sum_{\ell \in L: \ell < \ell^*} \ell \leq \ell^*$ again because the terms of L follow a geometric progression.

Finally, if $\ell(v) > \ell^*$ then each individual item can contribute at most $\frac{k}{\ell(v)} + 1$ to the revenue, because this is the buyer's per-item value. And by the definition of ℓ^* there can be at most $t\ell^*$ items such that $p_j \leq \frac{k}{\ell(v)} + 1 \leq \frac{k}{\ell^*} + 1$; only these items may generate revenue. Thus, we have

$$\sum_{\ell \in L: \ell > \ell^*} \mathbb{E}_v [\text{REV}_i(v, p, S) \mid \mathcal{E}_i^B(v, p, S) \wedge \ell(v) = \ell] \leq t\ell^* \cdot \sum_{\ell \in L: \ell > \ell^*} \left(1 + \frac{k}{\ell}\right) \leq t\ell^* \cdot |L| + tk, \quad (16)$$

where the last inequality once again follows because the terms of L follow a geometric progression, so the sum of the ℓ^{-1} terms is at most 1.

Combining everything, applying (14), (15), and (16) to (13) we have

$$\begin{aligned} \mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p)] &= \frac{1}{|L|} \cdot \sum_{\ell \in L} \mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p) \wedge \ell(v) = \ell] \\ &\leq \frac{1}{|L|} \cdot (2t(k + \ell^*) + t\ell^*|L| + tk), \end{aligned}$$

and since $\ell^* \leq \max(L) = \sqrt{k}$ and $|L| = \frac{1}{2} \log k$, we have that

$$\mathbb{E}_{v_i} [\text{REV}(v_i, p) \mid \mathcal{E}^B(v_i, p)] \leq c_2 \cdot \frac{tk}{\log k}$$

for some constant c_2 and all suitably large k , as desired.

This holds in both cases, completing the proof. \square

Before proving the lemma 17, recall the definitions of an *assignment* σ of items to buyers and the definition for $I \subseteq [n]$ to be $(1-c)$ -feasible under σ (Definitions 4 and 5). We will use the following more general claim:

Lemma 24. *Let $\mathcal{E}(c, C, n)$ denote the event that there exist some subset of buyers $I \subseteq [n]$ and assignment $\sigma : [m] \rightarrow [n] \cup \{\perp\}$ such that $|I| \geq Cn$ and I is $(1-c)$ -feasible under σ . For any fixed constant $\gamma > 0$, if $c = \gamma \cdot \log^{-1/2} n$ and $C = 8c = 8\gamma \log^{-1/2} n$, then*

$$\Pr_{v \sim \mathcal{D}} [\mathcal{E}(c, C, n)] = O\left(e^{-n^2 / \log^2 n}\right).$$

Proof of Lemma 24. To begin, consider a fixed $I \subseteq [n]$ let $\alpha := |I|/n$. We will be interested in the I for which $\alpha \geq C$. Next let $N(I) := \cup_{i \in I} A_i$ be the union of the A -sets for $i \in I$, and let $A_i^\sigma := A_i \cap \sigma^{-1}(i)$, so that σ is $(1-c)$ -feasible for i precisely when $|A_i^\sigma| \geq (1-c)|A_i| = (1-c)n$.

Next observe that for all assignments σ ,

$$\min_{i \in I} |A_i^\sigma| \leq \frac{1}{|I|} \cdot \sum_{i \in I} |A_i^\sigma| \leq \frac{1}{|I|} \cdot |N(I)|,$$

since $A_i^\sigma \subseteq A_i$ by definition and the A_i^σ are disjoint because σ is a well-defined mapping. Therefore if σ is $(1-c)$ -feasible for I , then

$$|I|n(1-c) \leq |I| \min_{i \in I} |A_i^\sigma| \leq |N(I)|. \quad (17)$$

We will argue that for fixed I a $(1-c)$ -feasible assignment for I is unlikely to exist, because this would imply that $|N(I)| \geq |I|(1-c)n$ by (17), which we will show is unlikely.

To this end, let $X_{ij} := \mathbb{1}\{j \in A_i\}$ indicate the event that $j \in A_i$; then let $Y_j := \max_{i \in I} X_{ij}$ indicate that j is in at least one set A_i for $i \in I$. Then observe that $Y := \sum_j Y_j = |N(I)|$.

Claim 25. $\{Y_j\}$ is negatively associated (NA).

Proof. This follows from results about NA variables by Joag-Dev and Proschan [1983]; in particular closure and composition properties for collections of NA variables.

First, consider each collection of indicator random variables $\mathcal{X}_i := \{X_{i'j} : i' = i\}$. Exactly n of these are 1, indicating the subset A chosen uniformly at random from $[m]$; therefore this is a permutation distribution, and so by [Joag-Dev and Proschan, 1983] Theorem 2.11 this collection \mathcal{X}_i is NA. Next, consider the larger collection $\mathcal{X} := \{X_{ij}\}$; as the union of independent collections of NA variables, \mathcal{X} is also NA ([Joag-Dev and Proschan, 1983], Property P_7). Finally, consider the collection $\mathcal{Y} := \{Y_j\}$. Since $Y_j := \max_{i \in I} X_{ij}$, each Y_j is defined on a unique subset of the NA variables \mathcal{X} ; for any Y_j and $Y_{j'}$, the sets of X_{ij} upon which they depend are disjoint. Since $\max(\cdot)$ is an increasing function, therefore \mathcal{Y} is NA also ([Joag-Dev and Proschan, 1983], Property P_6). \square

Bounding $|N(I)|$: Since Claim 25 demonstrates that \mathcal{Y} is negatively associated, we can use concentration. First note that for all j ,

$$\mathbb{E}_{v \sim \mathcal{D}} [Y_j] = 1 - \Pr_{v \sim \mathcal{D}} \left[\bigwedge_{i \in I} (j \notin A_i) \right] = 1 - (1 - 1/n)^{|I|} = 1 - (1 - 1/n)^{\alpha n}.$$

Then we may bound the probability that $|N(I)| \geq |I|(1-c)n$ by applying a Hoeffding bound to $\sum_j Y_j$ for any I with $|I| \geq Cn$ as follows:

$$\begin{aligned} \Pr_{v \sim \mathcal{D}} [|N(I)| \geq \alpha n^2(1-c)] &= \Pr_{v \sim \mathcal{D}} \left[Y - \mathbb{E}_{v \sim \mathcal{D}} [Y] \geq \alpha n^2(1-c) - \mathbb{E}_{v \sim \mathcal{D}} [Y] \right] \\ &\leq \exp \left(\frac{-2 \left(\alpha n^2(1-c) - \mathbb{E}_{v \sim \mathcal{D}} [Y] \right)^2}{n^2} \right) \\ &= \exp \left(\frac{-2 \left(\alpha n^2(1-c) - n^2(1 - (1 - 1/n)^{\alpha n}) \right)^2}{n^2} \right). \end{aligned}$$

Since $1 - (1 - 1/n)^{\alpha n} \leq 1 - e^{-\alpha}(1 - O(1/n))^\alpha \leq 1 - e^{-\alpha} + e^{-\alpha} \cdot O(1/n) \leq \alpha - \alpha^2/2 + \alpha^3/6 + e^{-\alpha} \cdot O(1/n)$, therefore

$$\begin{aligned} \Pr_{v \sim \mathcal{D}} [|N(I)| \geq \alpha n^2(1-c)] &\leq \exp \left(-2\alpha^2 n^2 \left((1-c) - \left(1 - \frac{\alpha}{2} + \frac{\alpha^2}{6} + \frac{e^{-\alpha}}{\alpha} O(n^{-1})\right) \right)^2 \right) \\ &= \exp \left(-2\alpha^2 n^2 \left(\frac{\alpha}{2} - \frac{\alpha^2}{6} - \frac{e^{-\alpha}}{\alpha} O(n^{-1}) - c \right)^2 \right) \\ &\leq \exp \left(-\frac{1}{8} n^2 \alpha^4 \left(1 - \frac{2\alpha}{3} - 4 \frac{e^{-\alpha}}{\alpha^2} O(n^{-1}) \right)^2 \right), \end{aligned}$$

since $\alpha \geq C > 4c$, and since $\alpha \geq C \gg n^{-1/3}$ for sufficiently large n this last term is $o(\alpha)$, so

$$\begin{aligned} &\leq \exp \left(-\frac{1}{8} n^2 \alpha^4 \left(1 - \frac{3\alpha}{4} \right)^2 \right) \\ &\leq \exp \left(-\frac{9}{128} n^2 \alpha^4 \right), \end{aligned}$$

where we used that $1 - 3\alpha/4 \geq 3/4$. Since I is a set for which $\alpha \geq C$, by (17), we therefore know that

$$\Pr_{v \sim \mathcal{D}} [\exists \sigma : \sigma \text{ is } (1-c)\text{-feasible for } I] \leq \exp \left(-\frac{9}{128} n^2 C^4 \right). \quad (18)$$

Concluding: We finish our argument with a union bound over the sets $I \subseteq [n]$ for which $|I| \geq Cn$. There are at most $2^n < e^n$ such sets; therefore by (18),

$$\begin{aligned} \Pr_{v \sim \mathcal{D}} \left[\bigvee_{I \subseteq [n]: |I| \geq Cn} (\exists \sigma : \sigma \text{ is } (1-c)\text{-feasible for } I) \right] &\leq \sum_{I \subseteq [n]: |I| \geq Cn} \Pr_{v \sim \mathcal{D}} [\exists \sigma : \sigma \text{ is } (1-c)\text{-feasible for } I] \\ &\leq \exp \left(n - \frac{9}{128} \cdot n^2 C^4 \right) \\ &= O(e^{-n^2 / \log^2 n}) \end{aligned}$$

for our choice of $C = O(\log^{-1/2}(n))$. The left-hand side is precisely the probability of the event $\mathcal{E}(c, C, n)$, and so this demonstrates the stated claim. \square

With this lemma in hand, Lemma 17 follows as a direct corollary.

Lemma 17. *Let $I_{\max}(v)$ be the largest $I \subseteq [n]$ for which I is $(1 - \frac{1}{t+1})$ -feasible under σ , for any assignment σ . Then for some c_3 and sufficiently large k , $\mathbb{E}_v [|I_{\max}(v)|] \leq c_3 \cdot \frac{n}{t}$.*

Proof. Since $t = \sqrt{\log n}$ for our choice of parameters, observe that any I which is $(1 - \frac{1}{t+1})$ -feasible under some assignment σ is also $(1 - \frac{2}{t})$ -feasible under σ .

We may therefore take $\gamma = 2$ and $c = \frac{2}{\log^{1/2} n}$ and $C = \frac{16}{\log^{1/2} n}$ and apply Lemma 24, which implies that

$$\Pr_{v \sim \mathcal{D}} [\mathcal{E}(c, C, n)] = O \left(e^{-n^2 / \log^2 n} \right).$$

Therefore, turning to $|I_{\max}(v)|$, we have

$$\begin{aligned} \mathbb{E}_v [|I_{\max}(v)|] &= \mathbb{E}_v [|I_{\max}(v)| \mid \neg \mathcal{E}(c, C, n)] \Pr_v [\neg \mathcal{E}(c, C, n)] + \mathbb{E}_v [|I_{\max}(v)| \mid \mathcal{E}(c, C, n)] \Pr_v [\mathcal{E}(c, C, n)] \\ &\leq Cn + n \cdot \Pr_v [\mathcal{E}(c, C, n)] \\ &\leq n \cdot \frac{16}{t} + n \cdot c_1 \cdot e^{-n^2 / \log^2 n} \leq 17 \cdot \frac{n}{t} \end{aligned}$$

for some constant c_1 and for sufficiently large $k = n$, as claimed. \square