



Buy-Many Mechanisms for Many Unit-Demand Buyers

Shuchi Chawla¹ , Rojin Rezvan¹ , Yifeng Teng² ,
and Christos Tzamos³

¹ University of Texas at Austin, Austin, TX 78712, USA
rojinrezvan@gmail.com

² Google Research, New York, NY, USA

³ University of Athens, Athens, Greece

Abstract. A recent line of research has established a novel desideratum for designing approximately-revenue-optimal multi-item mechanisms, namely the buy-many constraint. Under this constraint, prices for different allocations made by the mechanism must be subadditive, implying that the price of a bundle cannot exceed the sum of prices of individual items it contains. This natural constraint has enabled several positive results in multi-item mechanism design bypassing well-established impossibility results. Our work addresses the main open question from this literature of extending the buy-many constraint to multiple buyer settings and developing an approximation.

We propose a new revenue benchmark for multi-buyer mechanisms via an ex-ante relaxation that captures several different ways of extending the buy-many constraint to the multi-buyer setting. Our main result is that a simple sequential item pricing mechanism with buyer-specific prices can achieve an $O(\log m)$ approximation to this revenue benchmark when all buyers have unit-demand or additive preferences over m items. This is the best possible as it directly matches the previous results for the single-buyer setting where no simple mechanism can obtain a better approximation.

From a technical viewpoint we make two novel contributions. First, we develop a supply-constrained version of buy-many approximation for a single buyer. Second, we develop a multi-dimensional online contention resolution scheme for unit-demand buyers that may be of independent interest in mechanism design.

Keywords: Buy-many Mechanisms · Sequential Item Pricing · Multi-item Mechanism Design

1 Introduction

Revenue maximization in multi-parameter settings is notoriously challenging. It is known, for example, that in the absence of strong assumptions on the buyer's value distribution, the optimal revenue cannot be approximated within any finite factor by any mechanism with finite description complexity even for the simplest possible

setting of two items and a single unit-demand buyer [4, 17]. These impossibility results motivate the search for a different benchmark that captures salient features of the problem space while also permitting non-trivial approximation.

For single-agent settings one such benchmark was proposed by Briest *et al.* [4] and Chawla *et al.* [10]. These works showed that the infinite gap between the revenue of the optimal mechanism and any simple mechanism arises precisely when the mechanism offers “super-additive” options that charge for some bundles of items more than the sum of prices of their individual components. The so-called “buy-many” constraint disallows such exploitative behavior: when interpreted as a pricing over all possible allocations, the mechanism should be subadditive (defined appropriately over randomized allocations, as described in Sect. 2). An alternate motivation for this constraint arises as a consequence of buyer behavior in scenarios where the buyer can interact with the mechanism multiple times, purchasing an option from the menu each time. The buyer can then construct an allocation by purchasing a multiset of options from the menu in the cheapest possible manner. The buy-many constraint restricts the kinds of mechanisms the seller can use to extract revenue from the buyer. The corresponding benchmark is the optimal revenue that can be obtained from any mechanism satisfying the constraint. [4] and [10] showed that imposing such a constraint enables positive results even without requiring any assumptions on the buyer’s valuation function or the value distribution: for settings with a single buyer and m items, the optimal revenue obtained from any mechanism satisfying the buy-many constraint is no more than $O(\log m)$ times the revenue obtained from an item pricing.

Our work addresses the primary direction left open by these works and their followups, namely extending the buy-many constraint and revenue benchmark to settings with multiple buyers. Obtaining such an extension is challenging, however, as it depends on the particular implementation of the mechanism, differences in these details can lead to very different benchmarks. Indeed while the two approaches described above for formalizing the buy-many constraint – as a restriction on the pricing function and as a consequence of buyer behavior – lead to equivalent definitions in the single-buyer setting, they turn out to be very different in the multi-buyer setting. In fact, the latter approach of allowing buyers to interact with the mechanism multiple times provides no meaningful restriction on mechanisms at all and once again allows for unbounded revenue gaps between simple and optimal mechanisms.¹ We instead define the buy-many

¹ Consider, for example, multiple interactions of a buyer with the mechanism interleaved by purchases made by other buyers. As the item supply changes, the mechanism can update the prices on its menu, and no longer necessarily needs to satisfy a subadditivity constraint on the final pricing observed by the buyer. In fact, by exploiting this supply-based pricing approach, a multi-buyer buy-many mechanism can simulate any single-agent non-buy-many mechanism, inheriting the unbounded simple-versus-optimal revenue gaps of the latter setting. Sybil-proofness or false-name-proofness is even easier to achieve in principle, unless some symmetry-type restrictions are placed on the mechanism (as in [18, 25], for example), as the mechanism can simply refuse to make any allocations unless the number of agents is exactly n .

constraint as a restriction on the effective pricing faced by any individual buyer participating in the mechanism. We view the buy-many constraint as a standalone desideratum for the design of mechanisms that disallows exploitative behavior on part of the seller. This approach leads to a non-trivial restriction, while continuing to permit a range of different designs. We show that it becomes possible to recover upper bounds on the revenue gap between simple and optimal mechanisms without assumptions about buyers' value distributions, just as in the single-buyer setting.

Even as a restriction on the mechanism's pricing, the buy-many constraint can take many different forms depending on the information available to buyers at different stages in the mechanism. To obtain a comprehensive theory of multi-buyer buy-many mechanisms without going into the intricacies of specific applications, we avoid choosing any particular such extension, and instead define an ex-ante relaxation of the optimal buy-many revenue that simultaneously captures a broad range of settings. In Sect. 2, we describe the different forms of the buy-many constraint encompassed by this relaxation.

An Ex-Ante Relaxation of Buy-Many Revenue

We view multi-buyer settings as a collection of single-buyer instances via an ex-ante relaxation that allows the mechanism to allocate each item multiple times as long as the *expected* number of buyers each item is allocated to is at most 1. To be specific, let x_i be an m -dimensional allocation vector, with $x_{ij} \in [0, 1]$ denoting the probability of allocation of item j to a particular buyer i . We consider single-buyer mechanisms satisfying two restrictions. First, over the randomness in the buyer's valuation function, we require that the probability of allocation of each item j to the buyer is at most x_{ij} . We say that the mechanism satisfies the ex-ante constraint x_i . Second, we require that the mechanism satisfies the single-buyer buy-many constraint. We then consider the maximum revenue that can be obtained from any mechanism for buyer i that satisfies both of the aforementioned constraints. Note that the buy-many constraint is not closed over convex combinations, so distributions over buy-many menus can potentially obtain higher revenue than buy-many menus themselves. We accordingly consider the maximum revenue obtainable from individual buy-many mechanisms or distributions over buy-many mechanisms that satisfy the ex-ante constraint x_i . Let $\text{BUYMANYREV}_i(x_i)$ denote this upper bound. The following program then gives an upper bound on the revenue of multi-buyer buy-many mechanisms:

$$\text{EXANTE-BUYMANYREV} := \max_{x_1, \dots, x_n \geq \mathbf{0}} \sum_i \text{BUYMANYREV}_i(x_i) \text{ s.t. } \sum_i x_i \preceq \mathbf{1}.$$

Here " $\preceq$ " means pointwise dominance: for any two vectors $x, y \in \mathbb{R}^m$, $y \preceq x$ means $y_j \leq x_j$ for every $j \in [m]$. Our goal in this work is to design simple multi-buyer mechanisms that are clearly buy-many but at the same time are competitive with this new benchmark.

Approximating the Buy-Many Revenue by Item Pricings

Although buy-many mechanisms can, in general, be quite complicated, we show that when all buyers are unit-demand or additive, the optimal buy-many revenue (as well as the ex-ante upper-bound) can be approximated via a simple class of mechanisms, namely sequential item pricings. We obtain an $O(\log m)$ approximation where m is the number of items, matching within constant factors the approximation achieved by Chawla *et al.* [10] and Briest *et al.* [4] for single buyer settings.

Theorem 1. *For n independent buyers that are unit-demand or additive over m items with value distribution $\mathfrak{D}$,*

$$\text{EXANTE-BUYMANYREV}(\mathfrak{D}) \leq O(\log m) \text{SREV}(\mathfrak{D}).$$

Here SREV denotes the optimal revenue achievable through sequential item pricings. A sequential item pricing mechanism interacts with buyers sequentially in a particular order. It offers to each buyer the set of remaining items at pre-determined item prices and allows the buyer to purchase any item of her choice. Notably, our approximation result holds for a worst-case order of arrival of the buyers. Furthermore, the item prices offered to each buyer are non-adaptive in that prices are computed once at the beginning of the mechanism before buyer values are instantiated, and do not change based on instantiated values and purchasing decisions of buyers that arrive earlier in the ordering. Interestingly, sequential item pricing was previously shown in [7] to obtain a constant-factor approximation to the overall (non-buy-many) optimal revenue for unit-demand buyers when buyers' values are independent across items.

Our Techniques

Our main approximation result consists of two parts, each of which is of independent interest. First, we define an ex-ante supply-constrained relaxation of (distributions over) item pricings, similar to the ex-ante relaxation of general buy-many mechanisms discussed above. We show that the ex-ante item pricing revenue provides a logarithmic approximation to the ex-ante buy-many revenue for buyers with any combinatorial valuation function.

Focusing on ex-ante relaxations allows us to revert back to the single-agent setting. We extend the logarithmic upper bound on the gap between buy-many and item pricing revenues proved by [10] to the single buyer setting with a supply constraint. However, extending [10]'s argument is not straightforward because it does not provide any control on how the expected allocation of the item pricing relates to that of the optimal buy-many mechanism. We instead consider a Lagrangian version of the limited supply setting: in this setting, we are given a production cost for each item and our goal is to maximize the mechanism's profit, namely its revenue minus the expected cost it has to pay to produce the items it sells. Unfortunately, it turns out that for some cost vectors, the expected profit of a buy-many mechanism can be an $\Omega(m)$ factor larger than the expected

profit of any item pricing. We show instead that the optimal item pricing profit approximates the profit of an optimal buy-many mechanism that faces slightly higher (in multiplicative terms) production costs. This allows us to obtain a logarithmic bound on the gap between the two ex-ante relaxations.

Theorem 2. *For any joint value distribution $\mathfrak{D}$ over n buyers and m items,*

$$\text{EXANTE-BUYMANYREV}(\mathfrak{D}) \leq O(\log m) \text{EXANTE-SREV}(\mathfrak{D}).$$

A Multi-dimensional Contention Resolution Scheme. The second part of our argument relates the ex-ante item pricing revenue to the revenue of a non-adaptive sequential item pricing for unit-demand or additive buyers. This part employs an argument reminiscent of prophet inequalities and contention resolutions schemes (see, e.g., Feldman *et al.* [12]). Specifically, let x_i denote the allocation vector for agent i in the ex-ante-optimal item pricing revenue. For any given fixed ordering over the buyers, we construct pricings $\{q_i\}$ such that: (1) every buyer i obtains an expected allocation of at least $x_i/2$, and (2) the revenue obtained by the item pricing q_i from agent i is at least half the agent's contribution to the ex-ante item pricing revenue. Both of these properties are easy to observe for additive buyers, but challenging for unit-demand buyers. The challenge in showing (1) is that the set of items available when agent i arrives depends upon the instantiations of previous agents' values: as this set changes, the buyer's choice of what to buy also changes. (2) is challenging because revenue is not linear in allocation probability. Our key observation is that for unit-demand buyers, item pricing revenue exhibits concavity as a function of allocation probabilities: given any item pricing p with allocation probabilities x and any $\alpha < 1$, we can find another item pricing that allocates items with probabilities at most αx and obtains at least an α fraction of p 's revenue. We leave open the question of extending this type of multi-dimensional contention resolution scheme (and in particular, achieving the following theorem) to other valuation functions.

Theorem 3. *For any joint value distribution $\mathfrak{D}$ over n unit-demand or additive buyers and m items,*

$$\text{EXANTE-SREV}(\mathfrak{D}) \leq 2\text{SREV}(\mathfrak{D}).$$

Theorem 1 follows immediately from Theorems 2 and 3. We prove Theorem 2 in Sect. 3 and Theorem 3 in Sect. 4.

In Sect. 5, we mention some **motivating examples** regarding multi-buyer buy-many mechanisms, and address what are the main difficulties of extending our results to all valuation functions.

Further Related Work

The buy-many constraint was first proposed as an alternative to unconstrained revenue maximization by Briest *et al.* [4] in the context of a single unit-demand

buyer. Chawla *et al.* [10] extended the notion to arbitrary single buyer settings. Chawla *et al.* chawla2021pricing present improved approximations for buy-many mechanisms in single buyer settings where buyers' values satisfy an additional ordering property that includes, for example, the so-called FedEx setting. In a different direction, Chawla *et al.* chawla2020menu study the menu size complexity and revenue continuity of single buyer buy-many mechanisms. All of these works focus on single buyer problems.

There is extensive literature on approximating the optimal non-buy-many revenue under various constraints on the buyers' value distributions in both single-buyer and multi-buyer settings. Interestingly, many of these approximation results are achieved through sequential item pricing (or, in some cases, grand bundle pricing) mechanisms. For example, Chawla *et al.* [7] show that for multiple settings of matroid feasibility constraints, sequential item pricings give a constant approximation to the optimal non-buy-many revenue. In particular, for multiple unit-demand buyers with independent values over all items, sequential item pricing gives a 6.75-approximation to the optimal revenue, and this competitive ratio was improved to 4 by Alaei *et al.* [1] and generalized to a setting where each item has multiple units. For buyers with more general distribution, either a sequential pricing mechanism with personalized item prices or a sequential item pricing with anonymous prices and entry fee gives a constant approximation to the optimal revenue when there are multiple fractional subadditive buyers [5] and an $O(\log \log m)$ approximation when buyers are subadditive [5, 11]. Ma and Simchi-Levi [21] consider additive-valued buyers and a seller facing production costs, a setting that we revisit in our proof of Theorem 4, and achieve approximations using two-part tariffs. Also see [3, 6, 8, 17, 19, 20, 22, 24] for some other previous work on using simple mechanisms to approximate the optimal non-buy-many revenue. For sequential item pricings to be able to approximate the revenue of the optimal non-buy-many mechanisms, an important assumption is that each buyer should have independent item values, which we do not assume in this work.

Finally, ex-ante relaxations, first introduced to multi-buyer mechanism design by Alaei *et al.* [1] and Yan and Qiqi [23], have emerged as a powerful tool for simplifying multi-buyer mechanism design problems by breaking them up into their single-agent counterparts. For example, Chawla *et al.* [9] and Cai *et al.* [5] employ this approach for designing sequential mechanisms that approximate the optimal revenue for multiple agents with subadditive values. Such a technique is also used for analyzing the setting where buyers have non-linear utilities [2, 13, 15], and determining the revelation gap of the optimal mechanisms [14, 16].

2 Definitions

We study the multidimensional mechanism design problem where the seller has m heterogeneous items to sell to n buyers with independent value distributions, and aims to maximize the revenue. The buyer i 's value vector over items is specified by a distribution $\mathcal{D}_i$ over all value functions $v_i : 2^{[m]} \rightarrow \mathbb{R}^{\geq 0}$. Each

of these such functions v_i assigns a non-negative value to any subset of items $S \subseteq [m]$. Let $\mathfrak{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ denote the joint distribution over all values. Let $\Delta_m = [0, 1]^m$ be the set of all possible randomized allocations.

Unit-Demand and Additive Buyers. In this paper, we focus on settings where every buyer is unit-demand or where every buyer is additive valued. We say that a buyer is unit-demand over all items, if the buyer is only interested in purchasing one item, and her value for any set of items is solely determined by the item that is most valuable to her. In other words, for any set of items $S \subseteq [m]$, $v_i(S) = \max_{j \in S} v_i(\{j\})$. We say that a buyer is additive if for all $S \subseteq [m]$, $v_i(S) = \sum_{j \in S} v_i(\{j\})$. For ease of notation, let $v_i(\{j\}) = v_{ij}$. We make no assumptions on each buyer's value distribution $\mathcal{D}_i$ – the buyer's values for different items can be arbitrarily correlated.

Single-Buyer Mechanisms. By the taxation principle, any single-buyer mechanism is equivalent to a menu of options, and the buyer can select an outcome that maximizes her utility. Without loss of generality, we can assume that the menu assigns a price to *every* randomized allocation, a.k.a. lottery, $\lambda \in \Delta_m$, and can therefore simply represent the mechanism by a pricing function p . $p(\lambda)$ denotes the price of the lottery λ . For a buyer of type v_i , her (expected) value for a lottery λ is defined by $v_i(\lambda) := \sum_j v_{ij} \lambda_j$, and her utility for the lottery is defined to be $u_{v_i, p}(\lambda) := v_i(\lambda) - p(\lambda)$. When p is clear from the context, we drop the subscript and write the buyer's utility as $u_{v_i}(\lambda)$. Henceforth, we will refer interchangeably to mechanisms as pricing functions.

Given a pricing function p , a buyer of type v chooses the utility-maximizing lottery, denoted $\lambda_{v, p} := \arg \max_{\lambda} u_{v, p}(\lambda)$. The buyer's utility in the mechanism p is given by $u_p(v) := v(\lambda_{v, p}) - p(\lambda_{v, p})$; the buyer's payment is $\text{REV}_p(v) := p(\lambda_{v, p})$. We write the revenue of the mechanism under buyer distribution $\mathcal{D}$ as $\text{REV}_p(\mathcal{D}) = \mathbb{E}_{v \sim \mathcal{D}} \text{REV}_p(v)$. The mechanism is called a buy-one mechanism since the buyer only purchases one option from the menu.

Single-Buyer Buy-Many Mechanisms. In single buyer settings, if the buyer is allowed to purchase multiple options from a mechanism's menu, we call the mechanism buy-many. In particular, in a buy-many mechanism the buyer can purchase a (random) sequence of lotteries, where each lottery in the sequence can depend adaptively on the instantiations of previous lotteries. At the end of the process, the buyer gets the union of all allocated items in each step and pays the sum of the prices of all purchased lotteries.

Any buy-many mechanism can be described by a buy-one pricing function that satisfies a buy-many constraint. Intuitively, a buy-one pricing function satisfies the buy-many constraint if the buyer always prefers to purchase a single option from the menu, even if she has the option to adaptively interact with the mechanism multiple times.² For example, for a pricing function p , let $\hat{p}_i$ denote the minimum cost of acquiring item i by repeatedly purchasing some lottery until the item is instantiated. Then, the buy-many constraint implies that for any λ ,

² For formal definitions, see [10].

$p(\lambda) \leq \hat{p} \cdot \lambda$. Pricing functions that satisfy the buy-many constraint are called buy-many pricings, and we use BUYMANY to denote the set of such functions.

Let $\text{BUYMANYREV}(\mathcal{D}) := \max_{p \in \text{BUYMANY}} \text{REV}_p(\mathcal{D})$ denote the revenue of the optimal buy-many mechanism for a buyer with value distribution $\mathcal{D}$.

Item Pricings and Sequential Item Pricings. In a single-buyer setting, a (deterministic) item pricing mechanism sets an item price $p_j \in \mathbb{R}_+$ for every item j ; a buyer with value function v_i purchases the item j that maximizes $v_{ij} - p_j$. We use the price vector $p = (p_1, \dots, p_m)$ to denote the item pricing. Denote by $\text{SREV}(\mathcal{D})$ the optimal revenue obtained by any item pricing for a buyer with distribution $\mathcal{D}$. We also define random item pricing mechanisms as distributions over deterministic item pricings; we continue to use p to denote such pricings, although p is now a random variable.

A sequential item pricing for a multi-buyer setting specifies a serving order σ being a permutation over the n buyers, and n item pricings $p_1, \dots, p_n$, such that at step i the seller posts the pricing $p_{\sigma(i)}$ to buyer $\sigma(i)$, and the buyer purchases her favorite among the items that are still available. We use $\text{SREV}(\mathfrak{D})$ to denote the optimal revenue obtained by a sequential item pricing mechanism for buyers with values drawn from the joint distribution $\mathfrak{D}$.

Profit Maximization Under Production Costs. For a single-buyer setting, we study the profit-maximizing problem of the seller where each item has a production cost. Let $c = (c_1, c_2, \dots, c_m) \in \mathbb{R}_+^m$ denote the vector of production costs. Then the profit of the single-buyer mechanism given by pricing p is defined to be the revenue minus the production costs of the items:

$$\text{PROFIT}_{p,c}(\mathcal{D}) = \mathbb{E}_{v \sim \mathcal{D}}[p(\lambda_{v,p}) - \lambda_{v,p} \cdot c].$$

We denote by $\text{SPROFIT}_c(\mathcal{D}) = \max_{\text{item pricing } p} \text{PROFIT}_{p,c}(\mathcal{D})$ the optimal profit achievable by item pricings for the buyer with distribution $\mathcal{D}$ when there are production costs c for all items. Similarly define

$$\text{BUYMANYPROFIT}_c(\mathcal{D}) = \max_{\text{buy-many } p} \text{PROFIT}_{p,c}(\mathcal{D})$$

to be the optimal profit achievable by buy-many mechanisms for the buyer with distribution $\mathcal{D}$ for the setting with production costs c .

Ex-Ante Constrained Revenue

Ex-ante relaxations are a powerful technique for reducing multi-buyer mechanism design problems to their single-buyer counterparts. The key idea is to relax the ex-post supply constraint on items to an ex-ante feasibility constraint that requires each item to be sold at most once in expectation.

Recall that x_{ij} denotes the probability of allocating item j to buyer i ; $x_i = (x_{i1}, \dots, x_{im})$ denotes the vector of allocations of all items to buyer i ; and $x = (x_1, \dots, x_n)$ denote the vector of all allocations.

We say that a single-buyer pricing function p_i satisfies the ex-ante constraint x_i with respect to the value distribution $\mathcal{D}_i$ if for all $j \in [m]$, the expected allocation of item j to the buyer is no more than x_{ij} : $\mathbb{E}_{v \sim \mathcal{D}_i}[\lambda_{v,p_i}] \leq x_i$. Note that this definition extends to random pricing functions p in the straightforward manner: we want the expected allocation to be bounded by x_i where the expectation is taken over both the buyer's value and the randomness in the mechanism. That is, $\mathbb{E}_{v \sim \mathcal{D}_i, p}[\lambda_{v,p}] \leq x_i$.

For a single buyer i and an ex-ante constraint x_i , we can now define the optimal revenue that can be obtained from the buyer subject to an ex-ante constraint for various classes of mechanisms. In particular, let Δ_{IP} denote the space of all distributions over item pricings, and Δ_{BM} denote the space of all distributions over buy-many pricings. Then we define:

$$\begin{aligned} \text{SREV}(\mathcal{D}_i, x_i) &= \max_{p \in \Delta_{IP}: p \text{ satisfies } x_i \text{ w.r.t. } \mathcal{D}_i} \text{REV}_p(\mathcal{D}_i), \quad \text{and,} \\ \text{BUYMANYREV}(\mathcal{D}_i, x_i) &= \max_{p \in \Delta_{BM}: p \text{ satisfies } x_i \text{ w.r.t. } \mathcal{D}_i} \text{REV}_p(\mathcal{D}_i). \end{aligned}$$

Given a combined vector x of ex-ante constraints for every buyer $i \in [m]$, we can write the collective revenue of the optimal single-buyer mechanisms that satisfy these constraints as:

$$\begin{aligned} \text{EA-SREV}(\mathfrak{D}, x) &= \sum_i \text{SREV}(\mathcal{D}_i, x_i), \quad \text{and,} \\ \text{EA-BUYMANYREV}(\mathfrak{D}, x) &= \sum_i \text{BUYMANYREV}(\mathcal{D}_i, x_i), \end{aligned}$$

Finally, we can define the ex-ante relaxation for each class of mechanisms.

$$\begin{aligned} \text{EA-SREV}(\mathfrak{D}) &= \max_{x: \sum_i x_{ij} \leq 1 \forall j \in [m]} \text{EA-SREV}(\mathfrak{D}, x), \quad \text{and,} \\ \text{EA-BUYMANYREV}(\mathfrak{D}) &= \max_{x: \sum_i x_{ij} \leq 1 \forall j \in [m]} \text{EA-BUYMANYREV}(\mathfrak{D}, x), \end{aligned}$$

Some Settings Captured by the Ex-Ante Relaxation

Our ex-ante relaxation captures approaches that define the buy-many constraint as a restriction on the effective pricing faced by any individual buyer participating in the mechanism. We now describe some specific such settings. We start with one of the simplest settings, in which buyers arrive one after the other and each buyer faces a single-buyer mechanism.

Sequential Buy-Many Mechanisms. Sequential mechanisms make offers to each buyer sequentially in a pre-specified order. The i -th buyer in the order is asked to choose which items to purchase among the subset of items that remain after buyers 1 through $i - 1$ have made their choices. A natural definition for the buy-many constraint for this multi-buyer setting boils down to offering a buy-many constrained mechanism to each buyer for any subset of items remaining,

i.e. the full-mechanism can be defined as a collection of single-buyer mechanisms $\mathcal{M}_{i,S} \in \text{BUYMANY}$ for any buyer i and any subset S of remaining items.

Ex-Post Buy-Many Mechanisms. A more flexible design space for multi-buyer buy-many mechanisms is to consider direct mechanisms in which buyers truthfully submit their valuations and conditional on the valuations of other buyers each buyer is faced with a single buyer buy-many mechanism. Buy-many mechanisms for this setting are specified as a collection of single buyer buy-many mechanisms $\mathcal{M}_{i,\vec{v}_{-i}} \in \text{BUYMANY}$ for any buyer i and any combination of valuation functions v_{-i} for the other buyers.

Bayesian Buy-Many Mechanisms. Another potential definition of multi-buyer buy-many mechanisms is to consider Bayesian settings in which we require the options any single buyer faces to be buy-many in expectation over the valuations of other buyers. If $q_i(v_i, \vec{v}_{-i})$ and $p_i(v_i, \vec{v}_{-i})$ is the allocation and price offered to buyer i when her value is v_i and the other buyers have valuations $\vec{v}_{-i}$, we would require that the single buyer mechanism with allocation probabilities $q_i(v_i) \triangleq \mathbb{E}_{\vec{v}_{-i}} q_i(v_i, \vec{v}_{-i})$ and price $p_i(v_i) \triangleq \mathbb{E}_{\vec{v}_{-i}} p_i(v_i, \vec{v}_{-i})$ to satisfy the buy-many constraint.

3 Relating the Ex-Ante Relaxations

In this section we bound the gap between the ex-ante optimal buy many revenue and the ex-ante optimal item pricing revenue when the seller faces many buyers. This is a generalization of single-buyer buy-many revenue approximations to supply constrained settings. We emphasize that for the results in this section we *do not* require any assumptions on the buyer's valuation function, such as that it is unit-demand or additive.

A note on notation: since we consider a single-buyer problem in this section, we drop the subscript i from most notation and simply denote buyer i 's valuation function as v ; her allocation vector as x ; the probability that item j is allocated to the buyer as x_j ; etc. We will write the ex-ante buy many and item pricing revenues of this buyer simply as $\text{BUYMANYREV}(\mathcal{D}, x)$ and $\text{SREV}(\mathcal{D}, x)$ respectively.

Theorem 4. *For any single buyer with value distribution $\mathcal{D}$ over m items and any ex-ante supply constraint $x \in \Delta_m$,*

$$\text{EA-BUYMANYREV}(\mathcal{D}, x) \leq O(\log m) \text{EA-SREV}(\mathcal{D}, x).$$

Applying this theorem to each of n buyers, we obtain the following corollary.

Theorem 2. *For any joint value distribution $\mathfrak{D}$ over n buyers and m items,*

$$\text{EA-BUYMANYREV}(\mathfrak{D}) \leq O(\log m) \text{EA-SREV}(\mathfrak{D}).$$

Before we present a formal proof of Theorem 4, let us describe the main ideas. Theorem 4 is a generalization of Theorem 1.1 from [10], which shows that the ratio of BUYMANYREV and SREV is bounded by $O(\log m)$ in the absence of an ex-ante supply constraint. The proof technique of [10] does not directly lend itself to the ex-ante setting because it does not provide much control over the allocation probability of the random item pricing it produces. Indeed, the (random) item pricing it returns is independent of the buyer's value distribution, whereas the allocation probabilities (that are expectations over values drawn from the distribution) necessarily depend on the value distribution. Instead of applying the approach of [10] directly, we first consider a Lagrangian version of the supply constrained problem.

To simplify the following discussion, we will hide the argument $\mathcal{D}$ from the respective revenue benchmarks. Viewing $\text{SREV}(x)$ and $\text{BUYMANYREV}(x)$ as two multi-variate functions over $x \in \Delta_m$, we first observe that for any fixed ex-ante constraint x^o , there exists a cost vector c^o such that x^o is the solution to the optimization problem $\max_x (\text{SREV}(x) - c^o \cdot x)$. Indeed, because $\text{SREV}(x)$ is concave³, $c^o = \nabla \text{SREV}(x^o)$ is such a function. Furthermore, c^o is a non-negative vector, and so the gap between $\text{SREV}(x^o)$ and $\text{BUYMANYREV}(x^o)$ is bounded by the gap between $\text{SREV}(x^o) - c^o \cdot x^o$ and $\text{BUYMANYREV}(x^o) - c^o \cdot x^o$. This motivates studying the Lagrangian problem of maximizing the profit of a mechanism subject to production costs c^o . In particular, for any value distribution $\mathcal{D}$, we have:

$$\max_x \frac{\text{EA-BUYMANYREV}(\mathcal{D}, x)}{\text{EA-SREV}(\mathcal{D}, x)} \leq \max_c \frac{\text{BUYMANYPROFIT}_c(\mathcal{D})}{\text{SPROFIT}_c(\mathcal{D})}$$

Unfortunately, the gap on the right hand side can be very large:

Theorem 5. *There exists a unit-demand value distribution $\mathcal{D}$ over m items and a cost vector $c \in \mathbb{R}^m$, such that $\text{BUYMANYPROFIT}_c = \Omega(m) \text{SPROFIT}_c(\mathcal{D})$.*

We instead provide a bi-criteria approximation for the Lagrangian problem. In particular, we compare the profit of the optimal item pricing for cost vector c with the profit of the optimal buy many mechanism with production costs $2c$. This suffices to imply Theorem 4 with a slight worsening in the approximation factor.

Theorem 6. *For any single buyer with value distribution $\mathcal{D}$ over m -items and production costs vector c ,*

$$\text{BUYMANYPROFIT}_{2c}(\mathcal{D}) \leq 2 \ln 4m \text{SPROFIT}_c(\mathcal{D}).$$

The rest of the section is organized as follows. We first show a complete proof of Theorem 4 based on Theorem 6. We then describe and verify the gap example in Theorem 5. Each of these components is self-contained.

³ In fact, the function $\text{REV}(x)$ defined as maximum revenue from any restricted set of mechanism with allocations at most x is concave. This is because for any two allocations x and y and coefficient $1 \geq \alpha \geq 0$, one can consider mechanisms $M(x)$ and $M(y)$ that define $\text{REV}(x)$ and $\text{REV}(y)$, and run the former with probability α and the latter with probability $(1 - \alpha)$. Then $\text{REV}(\alpha x + (1 - \alpha)y) \geq \alpha \text{REV}(x) + (1 - \alpha) \text{REV}(y)$.

3.1 Proof of Theorem 4

Proof (Proof of Theorem 4).

We first note that $\text{SREV}(x)$ is a concave function over x because it optimizes for revenue over random pricings. Fix an ex-ante constraint x^o and consider the function $g(x) := \text{SREV}(x) - x \cdot \nabla \text{SREV}(x^o)$. This function is maximized at $x = x^o$. Furthermore, since $\text{SREV}(x)$ is monotone non-decreasing, $\nabla \text{SREV}(x^o) \geq \mathbf{0}$, and so $c := \nabla \text{SREV}(x^o)$ can be thought of as a vector of production costs. Since $g(x)$ is exactly the profit of an item pricing for buyer distribution $\mathcal{D}$ with allocation x and production cost vector c for all items, we know that the random item pricing p that achieves $\text{SREV}(x^o)$ is also the optimal item pricing for a buyer with value distribution $\mathcal{D}$ and item production costs c , without any allocation constraint.⁴

Now consider the buy-many profit optimization problem with production costs $2c$. The optimal profit is given by $\text{BUYMANYPROFIT}_{2c}$. Restricting attention to buy many mechanisms that satisfy the ex-ante constraint x^o , we let $\text{BUYMANYPROFIT}_{2c}(x^o)$ denote the optimal profit obtained over that set of mechanisms. Then, we can apply Theorem 6 to obtain:

$$\begin{aligned}
 \text{SREV}(x^o) &= \text{SPROFIT}_c(x^o) + c \cdot x^o \\
 &= \text{SPROFIT}_c(\mathcal{D}) + c \cdot x^o \\
 &\geq \frac{1}{2 \ln 4m} \text{BUYMANYPROFIT}_{2c}(\mathcal{D}) + c \cdot x^o \\
 &\geq \frac{1}{2 \ln 4m} \text{BUYMANYPROFIT}_{2c}(x^o) + c \cdot x^o \\
 &= \frac{1}{2 \ln 4m} (\text{BUYMANYPROFIT}_{2c}(x^o) + 2c \cdot x^o) + c \cdot x^o \left(1 - \frac{1}{\ln 4m}\right) \\
 &\geq \frac{1}{2 \ln 4m} \text{BUYMANYREV}(x^o),
 \end{aligned}$$

Here the first line is true by extracting the terms of item costs; the second line is true since $x = x^o$ is optimal for profit under item costs c ; the third line is true by Theorem 6; the fourth line is true since adding an allocation restriction cannot increase profit; the last line is true since $\text{BUYMANYPROFIT}_{2c}(x^o) + 2c \cdot x^o \geq \text{BUYMANYREV}(x^o)$. This finishes the proof of the theorem.

3.2 Proof of Theorem 5

Proof (Proof of Theorem 5). Let $c = (0, 2^m, 2^m, \dots, 2^m)$. For every j such that $2 \leq j \leq m$, let $v^{(j)}$ be the following unit-demand value function: $v_1^{(j)} = 2^j$; $v_j^{(j)} = 2^m$; $v_k^{(j)} = 0$ for $k \notin \{1, j\}$. In other words, the buyer with value $v^{(j)}$ is only interested in two items 1 and j , with the value for the first item being 2^j , and that for item j being 2^m . Consider the following value distribution $\mathcal{D}$: for every j such that $2 \leq j \leq m$, with probability 2^{-j} , $v = v^{(j)}$; for the remaining probability, $v = \mathbf{0}$. Now we analyze SPROFIT_c and BUYMANYPROFIT_c .

⁴ For any buyer i , it is possible that p has allocation $y \leq x^o$, with $\text{SREV}(y) = \text{SREV}(x^o)$. However, for any item j such that $y_j < x_j^o$, since $\text{SREV}(y) = \text{SREV}(x^o)$, the gradient $\nabla \text{SREV}(x^o)$ has value $c_j = 0$ on the j th component. Thus the profit of item pricing p for the buyer with production costs c is still $\text{SREV}(x^o) - c \cdot x^o$, although the actual allocation is less than x^o .

For any item pricing p , consider its profit contribution from the first item and the rest of the items. Since the buyer's value for the first item forms a geometric distribution, the profit contribution from the first item is upper bounded by the revenue of selling only the first item, which is $O(1)$. For the rest of the items, since when the buyer purchases some item $j > 1$, the item must have a price at most $2^m = c_j$, this means that the profit contribution of item j is at most 0. Thus $\text{SPROFIT}_c = O(1)$.

Consider the following buy-many mechanism: for every $j \geq 2$, there is a menu entry with allocation $\lambda^{(j)}$ and price $p^{(j)} = 2^{j-1} + 2^{m-1}$, where $\lambda_1^{(j)} = \lambda_j^{(j)} = 0.5$. For any $\lambda \in \Delta_m$ that is not some $\lambda^{(j)}$, its price is determined by the cheapest way to adaptively purchase it with $\lambda^{(2)}, \dots, \lambda^{(m)}$. For every buyer of type $v^{(j)}$, she buys lottery $(\lambda^{(j)}, p^{(j)})$ in the mechanism with utility 0.⁵ Since item j is only of interest to buyer $v^{(j)}$, the buyer would not purchase any other set of lotteries in the mechanism. Since profit from buyer $v^{(j)}$ is 2^{j-1} , and she gets realized in $\mathcal{D}$ with probability 2^{-j} , the expected profit of the buy-many mechanism is $\Omega(m)$.

4 Approximation via Sequential Item Pricing

We will now focus on item pricings and prove Theorem 3. In particular, we show that non-adaptive sequential item pricings can obtain half of the ex-ante optimal item pricing revenue, regardless of the order in which buyers are served.

Theorem 3. (Restatement) *For any joint distribution $\mathfrak{D} = (\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n)$ over n unit-demand or additive buyers and any order σ on arrival of buyers, there exists a deterministic sequential item pricing q with buyers arriving in order σ , such that*

$$\text{REV}_{\mathcal{D}}(q) \geq \frac{1}{2} \text{EA-SREV}(\mathfrak{D}).$$

For additive buyers, the items impose no externalities on each other, and so the theorem follows immediately from the single-unit prophet inequality. Henceforth we focus on unit-demand buyers. Our argument is based loosely around online contention resolution schemes (OCRS) [12] and prophet inequality arguments. The idea is to start with the optimal solution to the ex ante item pricing revenue: $x^* := \arg \max_{x: \sum_i x_{ij} \leq 1 \forall j \in [m]} \sum_i \text{SREV}(\mathcal{D}_i, x_i)$. Then, given some ordering σ over the buyers, we try to mimic this allocation by choosing pricings for each buyer that ensure that the buyer receives allocation comparable to x_i^* . As in OCRS, we tradeoff assigning enough allocation to a buyer with maintaining a good probability that items remain available for future buyers.

A key difference in our setting relative to work on OCRS is that the latter mostly focuses on utilitarian objectives, e.g. social welfare, so that the tradeoff is easily quantified: choosing an alternative with half the probability of the ex ante optimum, for example, provides half its contribution to the objective. Chawla *et*

⁵ We can reduce the price of $\lambda^{(j)}$ by some small $\epsilon > 0$ to make each buyer type's utility be strictly positive.

al. [6] show how to apply this approach to revenue for unit demand buyers with values independent across items by transforming values to Myersonian virtual values. Unfortunately this approach does not extend to values correlated across items because in correlated settings it is not possible to assign virtual values to each individual value independent of other values.

We develop an alternate argument. For any single unit-demand buyer, we consider how the item pricing revenue changes as the allocation of the buyer is decreased from some intended allocation x^* to a new allocation y that is component-wise smaller. We show that if x^* is realized by item pricing p , then we can realize allocation y while obtaining revenue at least $y \cdot p$. In particular, uniformly scaling down allocations by some factor scales down revenue by no more than the same factor. This allows us to carry out the OCRS-style argument. We formalize the above claim as a lemma before providing a proof of Theorem 3.

In the following discussion, for any unit-demand buyer with value distribution $\mathcal{D}$, let $\mathbf{x}_{p,S}(\mathcal{D})$ be the allocation vector of item pricing p over the set of available items $S \subseteq [m]$. We remove the distribution $\mathcal{D}$ whenever it is clear from the context. When $S = [m]$, we use $\mathbf{x}_p$ instead of $\mathbf{x}_{p,[m]}$.

Lemma 1. *For any unit-demand buyer, any deterministic item pricing p , and any distribution over set S of available items, let $x^* = \mathbb{E}_S[\mathbf{x}_{p,S}]$ be the expected allocations of p conditioned on the available set of items being S . Then for any allocation vector $y \in \Delta_m$ such that $y \preceq x^*$, there exists a random item pricing q such that*

$$\mathbb{E}_{q,S}[\mathbf{x}_q(S)] = y, \text{ and } \text{REV}_{q,S} = y \cdot p.$$

Proof. We first prove the theorem when both S and p are deterministic. Then, we extend the proof to the case where the available set S can be possibly randomized.

For $p = (p_1, \dots, p_m) \in \mathbb{R}_+^m$ being a deterministic item pricing, assume that the set of available items is fixed to be some deterministic set S . For any set $T \subseteq S$, define item pricing p_T to be the pricing p restricted to items in T . In other words, $p_{T,j} = p_j$ for all $j \in T$, and $p_{T,j} = \infty$ otherwise. For ease of notation, let $x_T = \mathbf{x}_{p,T}$ be the allocation under available set T . Observe that for any $j \notin T$, $x_j^* = 0$ and for any $j \in T$, $x_T(j) \geq x_j^*$: the latter is true since under the same pricing, when fewer items are available, buyer types that purchase an item not in T may switch to purchase some item in T , while the other buyer types' incentives remain unchanged. Thus for $y \preceq x^*$, y is in the convex hull of the set of $2^{|S|}$ points $X = \{x_T | T \subseteq S\}$. Write $y = \sum_{T \subseteq S} \alpha_T x_T$ as a convex combination of vectors in X , here $\alpha_T \in [0, 1]$ for every T , and $\sum_T \alpha_T = 1$. Consider the following randomized item pricing q : with probability α_T , $q = p_T$, $\forall T \subseteq S$. Then the expected allocation of q is exactly y . On the other hand, the expected revenue is $\text{REV}_{q,S} = p \cdot y$ because whenever item j is sold in q , it is sold at a price of p_j .

Now let S be a random variable over sets of available items. By defining $x_T = \mathbb{E}_S \mathbf{x}_{p,T \cap S}$ to be the expected allocation of item pricing p under available item set $T \cap S$, the above proof still goes through. This finishes the proof of the lemma.

With the help of the lemma above, we are now ready to prove Theorem 3.

Proof (Proof of Theorem 3). For every buyer i , let p_i be the (randomized) item pricing defining EA-SREV($\mathcal{D}$), and x_i be the corresponding allocation vector. We may assume $\sigma_i = i$ without loss of generality. We build the desired deterministic sequential item pricing q incrementally for every buyer. At arrival of buyer i , we use the allocation vector x_i , the distribution over currently available items (where the randomness comes from previous buyers and item pricings) and Lemma 1 to produce a (randomized) item pricing vector q_i with the property that for S_i being the random set of available items to buyer i ,

$$\mathbb{E}_{q_i, S_i}[x_{q_i, S_i}(\mathcal{D}_i)] = \frac{x_i}{2}, \text{ and, } \mathbb{E}_{q_i, S_i}[\text{REV}_{q_i}(\mathcal{D}_i, S_i)] = \frac{x_i}{2} \cdot p_i = \frac{1}{2} \text{EA-SREV}(\mathcal{D}_i). \quad (1)$$

Here $\text{REV}_{q_i}(\mathcal{D}_i, S_i)$ denotes the revenue of item pricing q_i for buyer i conditioned on the available item set being S_i when i arrives. If such q_i exists for every i , then the random sequential item pricing q satisfies $\text{REV}_{\mathcal{D}}(q) \geq \frac{1}{2} \text{EA-SREV}(\mathcal{D})$. Thus there must exist a realization of q being a deterministic sequential item pricing satisfying the requirement of the theorem.

Now it suffices to show that item pricing q_i exists for (1), and the rest of the proof is dedicated to proving this. An important observation is that if by induction the item prices $q_{i'}$ satisfying (1) exist for every $i' < i$, then every item belongs to S_i with probability at least $\frac{1}{2}$. Such a observation is true by noticing that by union bound, the allocation of item j in the first $i - 1$ steps is at most $\sum_{i' < i} \frac{x_{i', j}}{2} \leq \sum_{i' \leq m} \frac{x_{i', j}}{2} \leq \frac{1}{2}$ since item j has a total ex-ante allocation at most 1.

When p_i is deterministic, since every element in $[m]$ exists with probability at least $\frac{1}{2}$ in S_i , we know that for any realized buyer type, her favorite item still remains with probability at least $\frac{1}{2}$ and she would not deviate to purchase something else. Thus $\mathbb{E}_{p_i, S_i}[\mathbf{x}_{p_i, S_i}] \succeq \frac{1}{2} x_i$. By Lemma 1, there exists a random item pricing q , such that $\mathbb{E}_{q, S_i}[\mathbf{x}_{q, S_i}] = \frac{1}{2} x_i$, and $\mathbb{E}_{q, S_i}[\text{REV}_{q, S_i}] = \frac{1}{2} x_i \cdot p_i = \frac{1}{2} \text{EA-SREV}(\mathcal{D}_i)$. Thus (1) is satisfied.

When p_i is random, consider any instantiation of p_i . The same as the reasoning in the previous paragraph, $\mathbb{E}_{p_i, S_i}[\mathbf{x}_{p_i, S_i}] \succeq \frac{1}{2} \mathbf{x}_{p_i}$ still holds, and there exists a random item pricing q_{p_i} such that $\mathbb{E}_{q_{p_i}, S_i}[\mathbf{x}_{q_{p_i}, S_i}] = \frac{1}{2} \mathbf{x}_{p_i}$, and $\mathbb{E}_{q_{p_i}, S_i}[\text{REV}_{q_{p_i}, S_i}] = \frac{1}{2} \mathbf{x}_{p_i} \cdot p_i$. Consider the following random item pricing q_i : firstly generate a realization of random item pricing p_i , then generate a realization of random item pricing q_{p_i} defined above. The expected allocation of q_i is

$$\mathbb{E}_{p_i} \mathbb{E}_{q_{p_i}, S}[\mathbf{x}_{q_{p_i}, S}] = \frac{1}{2} x_i,$$

while the expected revenue is

$$\mathbb{E}_{p_i} \mathbb{E}_{q_{p_i}, S}[\text{REV}_{q_{p_i}, S}] = \mathbb{E}_{p_i} \left[\frac{1}{2} \mathbf{x}_{p_i} \cdot p_i \right] = \frac{1}{2} \text{EA-SREV}(\mathcal{D}_i).$$

Thus q_i satisfies (1), which finishes the proof of the theorem.

5 Discussion

In this section, we mention some examples that motivate definition and further work in multi-buyer buy-many mechanisms. Moreover, we discuss why extension of our results to any valuation function is challenging.

Motivating Examples. Consider a seller who is selling multiple items in multiple markets, and faces a common supply constraint across these markets. The seller is free to choose a different selling mechanism in each market, but interacts with a buyer from each market in just one go. For example, imagine Amazon selling rare books in multiple markets (US, Europe, India). Within each market, Amazon will display a price schedule that does not change frequently based on purchase decisions in other markets (but updates availability). This price schedule could be different for different markets based on local preferences and demand. Within each market individually, given that prices will remain static over short periods of time, a buy-many constraint is a natural property to satisfy. This scenario fits directly within our model.

For another similar scenario consider a travel website like hotwire.com which offers deals on airline tickets, hotel rooms, etc., without revealing complete vendor information. In effect, it sells lotteries. This is another example with supply constraints where the mechanism may personalize prices for each potential buyer (e.g. based on which browser the buyer is using). If the seller uses a non-buy-many mechanism (e.g. if lotteries on multiple items are generally more expensive than individual prices on the items they contain) it would lose customers over time.

Difficulties of Extending the Results to All Valuations. One component of our argument, namely approximating the ex-ante buy many revenue by the ex-ante SRev holds for every possible value function. However, we don't know how to extend the second part of the argument - approximating the ex-ante SRev using sequential item pricing - for value functions that are not unit-demand or additive. This requires constructing a multi-dimensional prophet inequality. The key technical challenge for non-unit-demand valuations is in keeping track of and controlling how the probability that a particular subset of items is available to an agent depends on decisions of other buyers.

References

1. Alaei, S.: Bayesian combinatorial auctions: expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* **43**(2), 930–972 (2014)
2. Alaei, S., Fu, H., Haghpanah, N., Hartline, J.: The simple economics of approximately optimal auctions. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, pp. 628–637. IEEE (2013)
3. Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M.: A simple and approximately optimal mechanism for an additive buyer. In: 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, pp. 21–30. IEEE (2014)

4. Briest, P., Chawla, S., Kleinberg, R., Weinberg, S.M.: Pricing lotteries. *J. Econ. Theory* **156**, 144–174 (2015)
5. Cai, Y., Zhao, M.: Simple mechanisms for subadditive buyers via duality. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 170–183 (2017)
6. Chawla, S., Hartline, J.D., Kleinberg, R.: Algorithmic pricing via virtual valuations. In: *Proceedings of the 8th ACM Conference on Electronic Commerce*, pp. 243–251 (2007)
7. Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: *Proceedings of the forty-second ACM symposium on Theory of computing*, pp. 311–320 (2010)
8. Chawla, S., Malec, D., Sivan, B.: The power of randomness in Bayesian optimal mechanism design. *Games Econom. Behav.* **91**, 297–317 (2015)
9. Chawla, S., Miller, J.B.: Mechanism design for subadditive agents via an ex ante relaxation. In: *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 579–596 (2016)
10. Chawla, S., Teng, Y., Tzamos, C.: Buy-many mechanisms are not much better than item pricing. In: *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 237–238 (2019)
11. Dütting, P., Kesselheim, T., Lucier, B.: An $o(\log \log m)$ prophet inequality for subadditive combinatorial auctions. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 306–317 (2020). <https://doi.org/10.1109/FOCS46700.2020.00037>
12. Feldman, M., Svensson, O., Zenklusen, R.: Online contention resolution schemes. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1014–1033. SIAM (2016)
13. Feng, Y., Hartline, J., Li, Y.: Simple mechanisms for non-linear agents. *arXiv preprint arXiv:2003.00545* (2020)
14. Feng, Y., Hartline, J.D.: An end-to-end argument in mechanism design (prior-independent auctions for budgeted agents). In: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 404–415. IEEE (2018)
15. Feng, Y., Hartline, J.D., Li, Y.: Optimal auctions vs. anonymous pricing: beyond linear utility. In: *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 885–886 (2019)
16. Feng, Y., Hartline, J.D., Li, Y.: Revelation gap for pricing from samples. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1438–1451 (2021)
17. Hart, S., Nisan, N.: Selling multiple correlated goods: revenue maximization and menu-size complexity. *J. Econ. Theory* **183**, 991–1029 (2019)
18. Iwasaki, A., et al.: Worst-case efficiency ratio in false-name-proof combinatorial auction mechanisms. In: van der Hoek, W., Kaminka, G.A., Lépérance, Y., Luck, M., Sen, S. (eds.) *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Toronto, Canada, 10–14 May 2010, vol. 1–3. pp. 633–640. IFAAMAS (2010)
19. Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, pp. 123–136 (2012)
20. Li, X., Yao, A.C.C.: On revenue maximization for selling multiple independently distributed items. *Proc. Nat. Acad. Sci.* **110**(28), 11232–11237 (2013)

21. Ma, W., Simchi-Levi, D.: Reaping the benefits of bundling under high production costs. In: Banerjee, A., Fukumizu, K. (eds.) The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, 13–15 April 2021, Virtual Event. Proceedings of Machine Learning Research, vol. 130, pp. 1342–1350. PMLR (2021)
22. Rubinstein, A., Weinberg, S.M.: Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Trans. Econ. Comput. (TEAC)* **6**(3–4), 1–25 (2018)
23. Yan, Q.: Mechanism design via correlation gap. In: Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 710–719. SIAM (2011)
24. Yao, A.C.C.: An n-to-1 bidder reduction for multi-item auctions and its applications. In: Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 92–109. SIAM (2014)
25. Yokoo, M., Sakurai, Y., Matsubara, S.: The effect of false-name bids in combinatorial auctions: new fraud in internet auctions. *Games Econ. Behav.* **46**(1), 174–188 (2004)