2009 50th Annual IEEE Symposium on Foundations of Computer Science

Dynamic and Non-Uniform Pricing Strategies for Revenue Maximization

Tanmoy Chakraborty* Zhiyi Huang Sanjeev Khanna†
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA, USA
Email: {tanmoy, hzhiyi, sanjeev}@cis.upenn.edu

Abstract— We study the ITEM PRICING problem for revenue maximization in the limited supply setting, where a single seller with n distinct items caters to m buyers with *unknown* subadditive valuation functions who arrive in a sequence. The seller sets the prices on individual items. Each buyer buys a subset of yet unsold items that maximizes her utility. Our goal is to design pricing strategies that guarantee an expected revenue that is within a small multiplicative factor of the optimal social welfare – an upper bound on the maximum revenue that can be generated by any pricing mechanism.

Most earlier work has focused on the unlimited supply setting, where selling an item to a buyer does not affect the availability of the item to the future buyers. Recently, Balcan et. al. [4] studied the limited supply setting, giving a randomized pricing strategy that achieves a $2^{O(\sqrt{\log n \log \log n})}$-approximation; their strategy assigns a single price to all items (uniform pricing), and never changes it (static pricing). They also showed that no pricing strategy that is both static and uniform can give better than $2^{\Omega(\log^{1/4} n)}$-approximation.

Our first result is a strengthening of the lower bound on approximation achievable by static uniform pricing to $2^{\Omega(\sqrt{\log n})}$. We then design *dynamic uniform pricing strategies* (all items are identically priced but item prices can change over time), that achieves $O(\log^2 n)$-approximation, and also show a lower bound of $\Omega\left((\log n/\log\log n)^2\right)$ for this class of strategies. Our strategies are simple to implement, and in particular, one strategy is to smoothly decrease the price over time. We also design a *static non-uniform pricing strategy* (different items can have different prices but prices do not change over time), that give poly-logarithmic approximation in a more restricted setting with few buyers.

Thus in the limited supply setting, our results highlight a strong separation between the power of dynamic and non-uniform pricing strategies versus static uniform pricing strategy. To our knowledge, this is the first non-trivial analysis of dynamic and non-uniform pricing schemes for revenue maximization in a setting with multiple distinct items.

Keywords-item pricing; limited supply setting; revenue maximization.

1. INTRODUCTION

We consider the following ITEM PRICING problem. Consider a finite set of items owned by a single seller, who wishes to sell them to multiple prospective buyers. The seller can price each item individually, and the price of a set of items is simply the sum of the prices of the individual items in the set. The buyers arrive in a sequence, and each buyer has her own *valuation function* $v(S)$, defined on every subset S of items. We assume that the valuation functions to be *subadditive*, which means that $v(S) + v(T) \geq v(S \cup T)$ for any pair of subsets S, T of items. For some results, we shall assume the valuations to be XOS, that is, they can be expressed as the maximum of several additive functions.

If a buyer buys a subset S of items S, her *utility* is defined as her valuation $v(S)$ of the set minus the price of the set S. Moreover, we assume the *limited supply* setting where a buyer can buy only yet unsold items. We assume that every buyer is selfish and rational, and thus always buy a subset of items that maximizes her utility. The strategy used by the seller in choosing the prices of the items is allowed to be randomized, and is referred to as a *pricing strategy*. The revenue obtained by the seller is the sum of the amounts paid by each buyer, and our goal is to design pricing strategies that maximize the expected revenue of the seller. This problem is made difficult by the fact that the seller has no knowledge of the valuation functions of the buyers, apart from the promise that they are subadditive. This is, for instance, in contrast to the Bayesian mechanism designs for revenue maximization, which assume that the valuation functions come from a known prior distribution. Optimal mechanisms, such as that given by Myerson [18], exist under this knowledge.

Pricing Strategies: A *uniform* pricing strategy is one where at any point of time, all unsold items are assigned the same price. The seller may set prices on the items initially and never change them, so that cost of an (unsold) item is the same for every buyer. We call such a strategy to be a *static* pricing strategy. Static pricing is the most widely applied pricing scheme till date. Alternatively, a seller may set fresh prices on the arrival of each buyer (without knowing the buyer's valuation function) – we shall call this a *dynamic* pricing strategy. Dynamic strategies have become more widely applicable with the introduction of online stores, since it is quite easy for online stores to show different prices to different customers. However, a dynamic strategy in which the price of an item fluctuates a lot may not be desirable in some applications. So we introduce an interesting subclass of dynamic strategies, called *dynamic*

* Supported in part by NSF Award CCF-0635084.
† Supported in part by a Guggenheim Fellowship, an IBM Faculty Award, and by NSF Award CCF-0635084.

0272-5428/09 $26.00 © 2009 IEEE
DOI 10.1109/FOCS.2009.43

monotone pricing strategies, where the price of an item can only decrease with time.

Social Welfare: An allocation of items involves distributing the items among the buyers, and the *social welfare* of an allocation the sum of the buyers' valuations for the items received by each of them. We denote the maximum social welfare, achieved by any allocation, by OPT. We measure the performance of a pricing strategy as the ratio of the maximum social welfare against the smallest expected revenue of the strategy, for any adversarially chosen ordering of the buyers. (Some of our results, where it will be explicitly stated, shall consider expected revenue under the assumption that the order in which buyers arrive is uniformly random.) If this ratio is at most α on any instance (where α can depend on the size of the instance), we say that the strategy achieves an α-approximation. Note that the maximum social welfare is an upper bound on the revenue the seller can obtain under any circumstance. In fact, there exists simple instances with n items and a single buyer where the maximum social welfare is $\log n$, but the revenue can never exceed 1 for any pricing function [4]. Thus we are comparing the performance of our strategies against a bar that is significantly higher than the optimal strategy, and we can never hope to achieve anything better than a logarithmic approximation. Our general goal is to design pricing strategies that achieve polylogarithmic approximation.

Related Work: The ITEM PRICING problem is closely related to the extensive body of literature in combinatorial auctions [7], which is the setting as described above, except that the buyers need not be arriving in a sequence but instead may place simultaneous bids on the items. A lot of recent literature has focused on social welfare maximization. This includes efficient approximation algorithms for computing maximum social welfare given oracle access to the valuation functions (eg. [11]), as well as on efficiently computable mechanisms that maximize social welfare and are truthful (eg. [17], [16], [10], [9]). For the first problem, Feige [11] gave a constant approximation for subadditive buyers, while for the second problem, Dobzinski et. al. [10], [9] gave logarithmic approximation when buyers have XOS valuations and subadditive valuations respectively. The mechanism achieving this approximation is in fact a static uniform pricing strategy.

A fair amount of research has focused on algorithms and truthful mechanisms for revenue maximization as well, but it has mostly considered the *unlimited supply setting* [15], where unlimited number of copies of each item is available to the seller. So one buyer receiving an item does not stop another buyer receiving the same item. Thus, the order in which buyers arrive has no effect on the performance of the mechanism, and in fact, the buyers can be handled independently. Some research has been directed towards developing new truthful mechanisms that maximize revenue [3], [12], [13], while others have focused on designing strategies for item pricing that maximizes revenue. The item pricing problem has received special attention because it is and has been the most widely applied mechanism for a seller wishing to sell items to potential buyers. All the research has focused only on designing static strategies (eg. [14], [2], [1], [6], [8]), and moreover, some of them have restricted their attention to finding envy-free pricing, which implies that the buyers come simultaneously, and the pricing must ensure that two buyers does not seek the same item. Moreover, most of these works assume severely restricted classes of valuation functions. For example, [14], [2] assume that all buyers are single-minded bidders. Their strategies were not only static but also uniform. Unsurprisingly, finding envy-free pricing is hard [8], and their results do not extend to more general buyer valuations such as XOS or subadditive. In all this work, the performance of a strategy has been measured as the ratio of the maximum social welfare to the expected revenue obtained.

More recently, Balcan, Blum and Mansour [4] considered static uniform pricing strategies with the objective of revenue maximization, in the limited supply setting, with subadditive buyer valuations. In the unlimited supply setting, they designed a pricing strategy that achieve revenue which is logarithmic approximation to the maximum social welfare even for general valuations. The strategy, again, was a uniform strategy. This result was also proved independently in [5]. However, in the limited supply setting, they could only get a $2^{O(\sqrt{\log n \log\log n})}$ factor approximation using a static uniform strategy. Crucially, they ruled out the existence of static uniform strategies that achieve anything better than a $2^{\Omega((\log n)^{1/4})}$ approximation, even if the buyer valuations are XOS, and the ordering of buyers is assumed to be chosen uniformly at random. Thus their result distinguished the limited and unlimited supply settings. This impossibility of getting a good (polylogarithmic) approximation is a consequence of being restricted to static uniform strategies, and it remains impossible even if the seller knew the buyer valuations, and had unlimited computational power. Further, almost all mechanisms in these related problems have only used a single price for all items. It is, therefore, natural to consider dropping one of these restrictions, namely, look at dynamic uniform strategies and static non-uniform strategies, both of which use multiple prices, and attempt to find better guarantees on the revenue.

Out Results and Techniques: The table below summarizes our results on the ITEM PRICING problem in the limited supply setting, along with relevant earlier work. Our contributions are labeled with the relevant theorem numbers.

We strengthen the hardness result of Balcan et. al. [4] by constructing instances with XOS valuations where uniform pricing functions cannot achieve any better than a $2^{\Omega(\sqrt{\log n})}$ approximation, even if the seller knew the buyer valuations,

Type of Pricing Strategy	Subadditive buyer valuations		ℓ-XOS buyer valuations	
	Algorithm [a]	**Lower Bound** [b]	**Algorithm** [a]	**Lower Bound** [b]
Dynamic Uniform Pricing	$O(\log^2 n)$ [Thm. 6] [c]	$\Omega\left(\left(\frac{\log n}{\log\log n}\right)^2\right)$ [Thm. 8]	$O(\log^2 n)$ [Thm. 6] [c]	$\Omega\left(\left(\frac{\log n}{\log\log n}\right)^2\right)$ [Thm. 8]
Dynamic Monotone Uniform Pricing	$O(\log^2 n)$ [Thm. 9] [d e]	$\Omega\left(\left(\frac{\log n}{\log\log n}\right)^2\right)$ [Thm. 8]	$O(\log^2 n)$ [Thm. 9] [d e]	$\Omega\left(\left(\frac{\log n}{\log\log n}\right)^2\right)$ [Thm. 8]
Static Uniform Pricing	$2^{O(\sqrt{\log n \log\log n})}$ [BBM08 [4]] [c]	$2^{\Omega(\sqrt{\log n})}$ [Thm. 5]	$2^{O(\sqrt{\log n \log\log n})}$ [BBM08 [4]] [c]	$2^{\Omega(\sqrt{\log n})}$ [Thm. 5]
Static Non-uniform Pricing	$2^{O(\sqrt{\log n \log\log n})}$ [BBM08 [4]] [c]	$\Omega(\log n)$ [BBM08 [4]]	$O\left(m \log \ell \log^3 n\right)$ [Thm. 11] [c]	$\Omega(\log n)$ [BBM08 [4]]

[a] All algorithms assume that the seller knows OPT up to a constant factor. This assumption can be removed by worsening the approximation ratio by a factor of $\log \mathrm{OPT}(\log\log \mathrm{OPT})^2$.

[b] All lower bounds are in the full information setting, where the seller knows the buyers' valuations, the number and arrival order of buyers, has unbounded computational power, and can even force the arrival order of buyers!

[c] Buyers arrive in an adversarial order. Thus the algorithm satisfies the upper bound for any order of buyers, including the order that minimizes expected revenue.

[d] Buyers arrive in a uniform random order, that is, every permutation of buyers is equally likely. The bound is on the expected revenue under this assumption.

[e] This algorithm also assumes that the seller knows the number of buyers m up to a constant factor, and it is deterministic. The assumption can be removed by making it randomized, and worsening the approximation ratio by a factor of $\log m(\log\log m)^2$.

and had unlimited computational power. Our basic construction is essentially similar to one given in [4]. We further extend our construction so that all buyers have the same XOS valuation function, so that the revenue is small for any order of buyers. Alternatively, we can extend it so that every buyer has an XOS buyer valuation (possibly different from the other buyers) that can be expressed as the maximum of three additive functions.

In contrast, we design a simple randomized dynamic uniform pricing strategy such that its expected revenue is $O(\log^2 n)$ approximation of the optimal social welfare, when the valuation functions are subadditive. The strategy is to randomly choose a threshold at the beginning, and then in each round, randomly choose a price above the threshold (but less than OPT) and put this price on each unsold item. By using a fresh random price (from a suitable set of prices) in each round, we guarantee, in expectation, to collect a large fraction of the revenue that can be obtained in that round from the remaining items. A suitable choice of threshold guarantees that no buyer picks up too many items because they are cheap.

The dynamic uniform pricing strategy described above achieves a high revenue, but requires random fluctuation in the price of an unsold item. This may not be a desirable property in some applications. We design a dynamic monotone uniform pricing strategy where the price of any unsold item decreases smoothly over time. We show that if the ordering of buyers is assumed to be uniformly random, that is, all permutations of buyers are equally likely, then the expected revenue is an $O(\log^2 n)$ approximation of the optimal social welfare. The strategy is in fact deterministic provided the seller knows estimates of OPT and m up to a constant factor. Deterministic strategies giving good approximation in such limited information settings are rare. We emphasize here that our lower bound for static uniform pricing holds for any ordering of buyers.

We show that the performance of our dynamic uniform pricing strategies are almost optimal among all dynamic uniform strategies, by showing that even if the seller knew the buyer valuations, had unlimited computational power, and could even force a particular ordering of the buyers, there exists instances with XOS valuations where the seller can achieve a revenue of at most $O(\mathrm{OPT}(\log\log n)^2/\log^2 n)$ if she is restricted to choosing a uniform dynamic pricing.

Finally, we give a static non-uniform strategy that gives an $O(m \log \ell \log^3 n)$-approximation if the buyers' valuations are XOS valuations that can be expressed as the maximum of ℓ additive components. Note that when the order of buyers is adversarial, the hard instance for static uniform pricing has only two buyers, and their valuation functions are the maximum of $o(\log n)$ additive functions components. In particular, our strategy gives polylogarithmic (in n) approximation when the number of buyers are small (polylogarithmic in n), and has XOS valuations which are the maximum of quasi-polynomial (in n) additive components. It is worth noting that our lower bound for dynamic uniform strategies also satisfies these properties.

2. Preliminaries

In the ITEM PRICING problem, we are given a single seller with a set I of n items that she wishes to sell. There are m buyers, each with their own *valuation function* defined on all subsets of I. A buyer with valuation function v values a subset of items $S \subseteq I$ at $v(S)$. The buyers arrive in a sequence, and each buyer visits the seller exactly once. The seller is allowed to set a price on each item, and the price of a subset of items is the sum of the prices of items in that subset. For every item sold to the buyers, the seller receives the price of that item. Note that an item can be sold at most once. So a seller can only offer those items to a buyer that has not been sold to any previous buyer. The *revenue* obtained by the seller is the sum of the prices of all the sold items.

Each buyer buys a subset of the items shown to her that maximizes her *utility*, which is defined as the value of the subset minus the price of the subset. This is clearly the behavior that is most beneficial to the buyer. The ITEM PRICING problem is to design (possibly randomized) pricing strategies for the seller that maximizes the expected revenue of the seller.

Unless noted otherwise, all our algorithmic results will assume that the seller has no knowledge of the order of arrival of the buyers, total number of buyers, or the valuation functions of buyers. We refer to a setting as the *full-information setting* if all these parameters are known to the seller.

Valuation Functions: Throughout this paper, we will assume that the buyer valuation function v is *subadditive*, which means that $v(S)+v(T) \geq v(S \cup T)\ \forall S \subseteq I, T \subseteq I$. *Unless explicitly stated otherwise, this will be the only assumption on the buyer valuation functions.* For some results, we shall assume the buyer valuations to be more restrictive than subadditive.

Definition 1: A subadditive valuation function v is called an *XOS valuation* if it can be expressed as $v(S) = \max\{a_1(S), a_2(S) \dots a_\ell(S)\}\ \forall S \subseteq I$ on all subsets of items S, where $a_1, a_2 \dots a_\ell$ are non-negative additive functions. The functions $a_1, a_2, ..., a_\ell$ are referred to as the *additive valuation components* of the XOS valuation v. We say that v is an ℓ-XOS function if it can be expressed using at most ℓ additive valuation components.

We note that a 1-XOS function is simply an additive function, that all XOS valuations are subadditive, and that not all subadditive valuations can be expressed as XOS valuations.

Pricing Strategies: We will study the power of some natural classes of pricing strategies.

Definition 2: A pricing strategy is said to be *static* if the seller initially sets prices on all items, and never changes the prices in the future. A pricing strategy is said to be *dynamic* if the seller is allowed to change prices at any point in time. A dynamic pricing strategy is also said to be *monotone* if the price of every item is non-increasing over time.

Definition 3: A pricing strategy is said to be *uniform* if at all points in time, all unsold items are assigned the same price.

2.1. Notation

For a buyer B with a valuation function v, we use $\Phi(B, S, p)$ to denote a set of items that the buyer B may buy when presented with set S of items, each of which are priced at p. Since $v(S) - p|S|$ is the utility if the buyer buys the set S, so $\Phi(B, S, p) = \text{argmax}_{S \subseteq J} v(S) - p|S|$ maximizes the utility. Note that there may be multiple possible sets that maximize the utility. In this paper, when we make a statement involving $\Phi(B, S, p)$, the statement shall hold for any choice of these sets. We shall denote the maximum utility as $\mathcal{U}(B, S, p)$; note that in contrast to $\Phi(B, S, p)$, the value $\mathcal{U}(B, I, p)$ is uniquely defined. When the underlying buyer B is clear from the context, we shall denote these two values as $\Phi(S, p)$ and $\mathcal{U}(S, p)$ respectively. Moreover, if the set of available items S is also clear from the context, then we shall denote these two values as $\Phi(p)$ and $\mathcal{U}(p)$ respectively. For any set S and a buyer with valuation function v, we define $H_v(S) = \max_{S' \subseteq S} v(S')$ as the maximum utility the buyer can get if all items in S are offered to her at zero price.

Definition 4: We say that a set of items S is *supported at a price* p with respect to some buyer B with valuation function v, if B buys the entire set S when the set S is presented to B at a uniform pricing of p on each item.

The following lemma follows easily from the fact that the valuation functions are subadditive, and was proved by Balcan et. al. [4].

Lemma 1: Let S be a set of items that is supported at price p. with respect to a buyer B with valuation function v. Then $v(S') \geq p|S'|$ for all $S' \subseteq S$.

Proof: Suppose not. Then there exists $S' \subset S$ with $v(S') < p|S'|$. By subadditivity of v, we know $v(S') + v(S \setminus S') \geq v(S)$, and hence $v(S \setminus S') \geq v(S) - v(S')$. Then the utility for B of buying the set $(S \setminus S')$ is at least $v(S) - v(S') - p|S/S'|$. But

$$\begin{aligned}&(v(S) - v(S') - p|S/S'|) - p|S'| + p|S'| = \\ &\qquad (v(S) - p|S|) - v(S') + p|S'| > v(S) - p|S| \ ,\end{aligned}$$

contradicting the assumption in the lemma that buyer picks set S at price p. ■

2.2. Optimal Social Welfare and Revenue Approximation

We now define optimal social welfare, the measure against which we evaluate the performance of our pricing strategies.

Definition 5: An *allocation* of a set S of items to buyers $B_1, B_2 \dots B_m$ with valuations $v_1, v_2 \dots v_m$, respectively, is an m-tuple $(T_1, T_2, \dots, T_m)$ such that $T_i \subseteq S$ for $1 \leq i \leq m$, and $T_i \cap T_j = \emptyset$ for $1 \leq i, j \leq m$. The *social welfare* of an allocation is defined as $\sum_{i=1}^{m} v_i(T_i)$, and an allocation is said to be a *social welfare maximizing allocation* if it maximizes $\sum_{i=1}^{m} v_i(T_i)$. The *optimal social welfare* OPT is defined as the social welfare of a social welfare maximizing allocation.

Clearly, OPT is an upper bound on the revenue that any pricing strategy can get. Let R be the revenue obtained by the strategy, which is the sum of the amounts paid by all the buyers.

Definition 6: A pricing strategy is said to achieve an α-approximation if the expected revenue of the strategy $\mathbf{E}[R]$ is at least OPT/α.

Unless stated otherwise, the expected revenue is computed with adversarial ordering of the buyers, that is, the ordering that minimizes the expected revenue of the strategy. In other words, we require a strategy to work well irrespective of the order of buyers in which they arrive.

Note that OPT is not a tight upper bound on the maximum revenue that can be achieved by any pricing strategy, even with full knowledge of buyer valuations and unbounded computational power. In fact, the following example was given in Balcan et. al. [4]: if there is a single buyer with valuation function $v(S) = \sum_{i=1}^{|S|} 1/i$, then for any pricing of the n items, the revenue is at most 1, while $\mathrm{OPT} = \Theta(\log n)$. This shows that nothing better than a logarithmic approximation can be achieved in the absence of any other assumption on the buyer valuations.

2.3. The Single Buyer Setting with Uniform Pricing

Balcan et. al. [4] considered the setting where there is an unlimited supply of each item, so that no buyer is affected by items bought before her arrival. In particular, if there is only a single buyer, then there is no distinction between limited and unlimited supply, as long as the buyer never wants more than one copy of the same item. For the single buyer case, Balcan et. al. [4] gave an $O(\log n)$ approximation, and in the process proved some lemmas that will be useful for our algorithmic results in the limited supply setting as well.

Suppose a set S is being shown to a buyer B with valuation function v. The optimal social welfare in this single buyer instance is $H_v(S)$. We consider setting a uniform price, that is, the same price on all items. The following lemma states that the number of items bought monotonically decreases as the price on the items is increased. It was proved by Balcan et. al. [4].

Lemma 2: (Lemma 6 of [4]) Suppose a buyer B is offered a set S of items using a uniform pricing. Then for any $p > p' \geq 0$, if B buys $\Phi(p)$ if all items are priced at p, and $\Phi(p')$ if all items are priced at p', then $|\Phi(p)| \leq |\Phi(p')|$. Thus there exist prices $\infty = q_0 > q_1 > \ldots > q_l > q_{l+1} = 0$ and integers $0 = n_0 < n_1 < \ldots < n_l \leq |S|$ such that when items in S are uniformly priced at $p \in [q_{t+1}, q_t)$ there is a subset $S' \subseteq S$ of size n_t that is supported at price p, and the utility $\mathcal{U}(p)$ of the buyer B satisfies

$$\mathcal{U}(p) = \mathcal{U}(q_t) + n_t(q_t - p) \tag{1}$$

Since the empty set maximizes utility when the price is q_1, we get that $\mathcal{U}(q_1) = 0$. Moreover, the utility at price $q_{\ell+1} = 0$ is $\mathcal{U}(q_{\ell+1}) = H_v(S)$. Thus we get that $H_v(S) = \sum_{t=1}^{l} n_t(q_t - q_{t+1})$.

The following lemma is a slight variation of Lemma 8 of [4]. Its proof can be found in the full version.

Lemma 3: Suppose a set S is being shown to a buyer B, with valuation function v, using a uniform price. Let H' be any number such that $H' \geq H_v(S)$. Let $\gamma > 1$, and let $p[t] = H'/\gamma^t$. Then, for any $k \geq 0$, we have

$$\sum_{t=1}^{k} p[t]|\Phi(p[t])| \geq \frac{1}{\gamma - 1}\left(H_v(S) - \frac{|S|H'}{\gamma^k}\right) .$$

Briefly, Lemma 3 will be used as follows: if one of $\{H', H'/2, H'/4 \ldots H'/2^k\}$ is chosen uniformly at random and set as the uniform price for all items in S, then for a sufficiently large choice of k, the revenue obtained is $\Omega(H_v(S)/k)$. This will happen when the right-hand-side of the equation in the lemma evaluates to $\Omega(H_v(S))$. We shall frequently use this lemma, and with $H' = \Theta(H)$, our choice of k will be logarithmic in the number of items.

2.4. Optimizing with Unknown Parameters

Almost all our algorithms use the following lemma, which was implicitly mentioned in the Appendix of [4]. It tells us that strategies can be allowed to assume that it approximately knows the value of some parameters, as long as the parameters are not too large, since these assumptions can be removed by guessing the value of these parameters and getting it correct with inverse-polylogarithmic probability. The lemma below is applicable to the ITEM PRICING problem with multiple buyers, and to both static and dynamic pricing strategies. Its proof can be found in the full version.

Lemma 4: Consider a pricing strategy $\mathcal{S}$ that gives an α-approximation in expected revenue, provided the seller knows the value of some parameter x to within a factor of 2. Then if the seller instead only knows that $L \leq x < H$, there exists a pricing strategy $\mathcal{S}'$ that gives an $O(\alpha \log(H/L))$ approximation in expected revenue, where L and H are powers of 2. If the seller instead only knows that $x \geq 1$ but no upper bound, then for any constant $\epsilon > 0$, there exists a pricing strategy $\mathcal{S}''$ that gives an $O(\alpha \log x (\log\log x)^{1+\epsilon})$ approximation in expected revenue.

3. Improved Lower Bounds for Static Uniform Pricing

We present here improved lower bounds for static uniform pricing. The core of our construction is the same as the lower bound construction in Balcan et. al. [4], however, with appropriate modifications in the choice of parameters, our lower bound almost matches the upper bound in [4]. We are also able to strengthen our construction to the case of identical buyers as well as to the case where each buyer uses simple XOS functions with only 3 additive components.

The following theorem summarizes our lower bound results about static uniform pricing. We defer the proof to the full version of the paper.

Theorem 5: There exists a set of buyers with XOS valuations, such that if the seller is restricted to a static uniform pricing strategy, then even in the full information setting, for any choice of price, the revenue produced is at most $\mathrm{OPT}/2^{\Omega(\sqrt{\log n})}$, where n is the number of items. Additionally, one of the following (but not both) can also be ensured, with the revenue still being at most $\mathrm{OPT}/2^{\Omega(\sqrt{\log n})}$:

- The valuations of all the buyers can be expressed as 3-XOS functions.
- All buyers have identical valuation function.

4. Dynamic Uniform Pricing

We now present a dynamic uniform pricing strategy that achieves an $O(\log^2 n)$-approximation to the revenue when buyer valuations are subadditive. This improves upon the previous best known approximation factor of $2^{O(\sqrt{\log n \log\log n})}$ [4] for the Item Pricing problem. Our strategy makes the assumption that the seller knows OPT, the maximum social welfare, to within a constant factor. However, this assumption can easily be eliminated by using Lemma 4, worsening the approximation ratio of the strategy by a poly-logarithmic factor.

We will also establish an almost matching lower bound result which shows that no dynamic uniform pricing strategy can achieve $o(\log^2 n/\log\log^2 n)$-approximation even when buyers are restricted to XOS valuations, the seller knows the value of OPT, buyer valuation functions, and is allowed to specify the order of arrival of the buyers!

4.1. A Dynamic Uniform Pricing Algorithm

The algorithm follows a simple strategy. Let $k = \lceil \log n \rceil + 1$, and let $p_i = \mathrm{OPT}/2^i$ (recall that OPT denotes the maximum social welfare). The algorithm starts at time 0 by choosing a *threshold value* p^* from the set $\{p_1, p_2 \dots p_{k+1}\}$, uniformly at random. Upon arrival of any buyer, the algorithm chooses a price $\hat{p}$ uniformly at random from the set $\{p_1, p_2 \dots, p^*\}$, and assigns the price $\hat{p}$ to all items that are yet unsold.

Theorem 6: If the buyer valuations are subadditive, then the expected revenue obtained by the dynamic strategy above is $\Omega(\mathrm{OPT}/\log^2 n)$.

The following lemma is key to the proof of Theorem 6. It says that if the threshold is "correctly" chosen, then our dynamically resetting prices give a large fraction of the maximum possible revenue.

Lemma 7: Suppose when the i^{th} buyer B_i arrives, there remains a set L_i^j of unsold items such that $v_i(L_i^j) \geq p_j|L_i^j|$, where v_i is the valuation function of B_i. Then if the seller picks a price from $\{p_1, p_2, \cdots, p_{j+1}\}$ uniformly at random, and prices all items at this single price, it receives an expected revenue of at least $p_j|L_i^j|/2(j+1)$ from this buyer.

Proof: Let I' be the set of unsold items when the buyer B_i arrives. Since this is a single buyer setting with uniform pricing, Lemma 2 applies. Thus the number of items sold is a non-increasing function of the price set on all items, and equation (1) is applicable.

Now if the uniform price chosen by the seller is p_{j+1}, then buying the set L_i^j would give B_i a utility of at least $v_i(L_i^j) - p_{j+1}|L_i^j| \geq p_j|L_i^j| - p_{j+1}|L_i^j|$ since $v_i(L_i^j) \geq p_j|L_i^j|$ by the assumption of the lemma. Thus $\mathcal{U}(B_i, I', p_{j+1}) \geq p_j|L_i^j| - p_{j+1}|L_i^j| = p_j|L_i^j|/2$. We shall now find a suitable upper bound on $\mathcal{U}(B_i, I', p_{j+1})$.

Suppose $q_s > p_{j+1} \geq q_{s+1}$, for some $s \leq l$. Then, since $\mathcal{U}(B_i, I', q_0) = \mathcal{U}(B_i, I', q_1) = 0$, and also that $q_1 \leq \mathrm{OPT}$, we get that

$$\begin{aligned}
&\mathcal{U}(B_i, I', p_{j+1}) \\
&= \mathcal{U}(B_i, I', q_s) + (\mathcal{U}(B_i, I', p_{j+1}) - \mathcal{U}(B_i, I', q_s)) \\
&= \sum_{t=1}^{s-1} (\mathcal{U}(B_i, I', q_{t+1}) - \mathcal{U}(B_i, I', q_t)) + n_s(q_s - p_{j+1}) \\
&= \sum_{t=1}^{s-1} n_t(q_t - q_{t+1}) + n_s(q_s - p_{j+1}) \ .
\end{aligned}$$

The above sum can be seen as an integral of the following step function f from p_{j+1} to q_1: in the range $[q_{t+1}, q_t)$, f takes the value n_t. So we can upper bound it by an upper integral of f. Note that $f(p) \leq |\Phi(B_i, I', p)| \leq |S|$, and also that f is a decreasing function. Thus we get

$$\begin{aligned}
\mathcal{U}(B_i, I', p_{j+1}) &\leq \sum_{t=0}^{j} |\Phi(B_i, I', p_{t+1})|(p_t - p_{t+1}) \\
&= \sum_{t=1}^{j+1} |\Phi(B_i, I', p_t)| p_t \ .
\end{aligned}$$

So we get $\sum_{t=1}^{j+1} |\Phi(B_i, I, p_t)| p_t \geq \mathcal{U}(B_i, I', p_{j+1}) \geq p_j|L_i^j|/2$. Thus the expected revenue obtained from B_i is $\sum_{t=1}^{j+1} |\Phi(B_i, I, p_t)| p_t/(j+1) \geq p_j|L_i^j|/2(j+1)$, completing the proof of the lemma. ∎

Proof of Theorem 6: Let $(T_1, T_2, \dots T_m)$ be an optimal allocation of items to buyers $B_1, B_2 \dots B_m$, who has valuation functions $v_1, v_2 \dots v_m$ respectively, such that $\sum_{i=1}^{m} v_i(T_i) = \mathrm{OPT}$ is the maximum social welfare. Also, let T_i^j be the subset of T_i that would be bought by B_i if it were shown only the items in T_i, and all items were uniformly priced at p_j. Now consider the case when $p^* = p_{j+1}$. Let R^j be the revenue in this case. Let $Z_i^j \subseteq T_i^j$ be a random variable that denotes the subset of items in T_i^j that are sold before buyer B_i comes. Then $R^j \geq \sum_{i=1}^{m} p^*|Z_i^j| = \sum_{i=1}^{m} p_j|Z_i^j|/2$.

Note that $v_i(T_i^j \setminus Z_i^j) \geq p_j|T_i^j \setminus Z_i^j|$ by Lemma 1. So, by Lemma 7, conditioned on the set Z_i^j, the expected revenue received from B_i is at least $\left(p_j|T_i^j \setminus Z_i^j|\right)/2(j+1)$. Thus, conditioned on the sets Z_i^j for all i, we have

$$\begin{aligned} & \mathbf{E}[R^j | Z_i^j \ \forall 1 \le i \le m] \\ \ge \quad & \Omega\left(\sum_{i=1}^m \left(p_j|Z_i^j| + \frac{p_j|T_i^j \setminus Z_i^j|}{j}\right)\right) \\ = \quad & \Omega\left(\sum_{i=1}^m \frac{p_j|T_i^j|}{j}\right) \quad . \end{aligned}$$

Since the value on the right-hand side above is independent of the variables Z_i^j on which the expectation of R^j is conditioned on, we get that $\mathbf{E}[R^j] = \Omega\left(\sum_{i=1}^m p_j|T_i^j|/j\right)$.

Thus the expected revenue of our dynamic strategy is given by

$$\mathbf{E}[R] = \frac{1}{k+1}\sum_{j=0}^{k}\mathbf{E}[R^j] = \Omega\left(\sum_{j=0}^{k}\sum_{i=1}^{m}\frac{p_j|T_i^j|}{k^2}\right) \qquad (2)$$

Since $k = \lceil \log n \rceil$, and $\mathrm{OPT} \ge H_{v_i}(T_i)$, from Lemma 3 and Equation (2), it follows that

$$\sum_{j=0}^{k} p_j|T_i^j| \ge \Omega\left(v_i(T_i) - \frac{|T_i|\mathrm{OPT}}{2n}\right) \quad .$$

Thus we have

$$\begin{aligned} \mathbf{E}[R] &= \Omega\left(\frac{1}{k^2}\left(\sum_{i=1}^m v_i(T_i) - \sum_{i=1}^m \frac{|T_i|\mathrm{OPT}}{2n}\right)\right) \\ &= \Omega\left(\frac{1}{k^2}\left(\mathrm{OPT} - \frac{\mathrm{OPT}}{2}\right)\right) = \Omega\left(\frac{\mathrm{OPT}}{\log^2 n}\right) \quad . \end{aligned}$$

■

4.2. Lower Bound for Dynamic Uniform Pricing

We shall now construct a family of instances of the problem, where the buyers have distinct XOS valuations, with $O(\log n/\log\log n)$ additive components in each valuation function, such that no dynamic uniform strategy can achieve an $o(\log^2 n/\log\log^2 n)$-approximation, even in the full information setting, and when the seller can even specify the order in which the buyers should arrive.

Theorem 8: There exists a set of buyers with XOS valuations, such that if the seller is restricted to using a dynamic uniform pricing strategy, then even when the seller has full information of buyer valuation functions and can even choose the order of arrival of the buyers, the revenue produced is $O((\log\log n)^2/\log^2 n)$ times OPT, where n is the number of items.

Proof: Let $B_1, B_2 \dots B_m$ denote the buyers. Our construction will use three integer parameters k, F, and Y, to be specified later. These parameters will satisfy the conditions that $k > 1$, $F > 1$, $Y > 4$, and $m \ge 2Y \ge 4k$. Let $f(i) = (i+1)F/Y^i$. Then, $f(0) > f(1) > \dots > f(k) > f(k+1)$.

For each buyer B_i, we create $2(k+1)$ disjoint sets of items $S_{i0}, S_{i1} \dots S_{ik}$ and $S'_{i0}, S'_{i1} \dots S'_{ik}$ such that $|S_{ij}| = |S'_{ij}| = Y^j$ items each. Let $S_i = \cup_{0\le j\le k} S_{ij}$ and $S'_i = \cup_{0\le j\le k} S'_{ij}$. We call the items in S_i as *shared* and those in S'_i as *private*. The private items of B_i are valued by buyer i only, and has zero value to all other buyers.

The valuation function v_i of buyer B_i is defined as an XOS valuation with $(k+2)$ additive functions $v_{i0}, v_{i1} \dots v_{i(k+1)}$ in its support, that is, $v_i = \max_{0\le j\le k+1} v_{ij}$. Specifically, for $0 \le j \le k$, the valuation function v_{ij} has positive value only for private items, and

$$v_{ij}(x) = \begin{cases} f(j) & \text{if } x \in S'_{ij} \\ 0 & \text{otherwise} \end{cases}$$

The valuation function $v_{i(k+1)}$ has positive values only for shared items:

$$v_{i(k+1)}(x) = \begin{cases} f(j) & \text{if } x \in S_{ij},\ 0 \le j \le k \\ f(j+1) & \text{if } x \in S_{\ell j},\ 1 \le \ell \le m, \ell \ne i, \\ & \text{and } 0 \le j \le k \end{cases}$$

This completes the description of the instance. Note that $v_i(S'_{ij}) = f(j)|S'_{ij}| = (j+1)F$, and that $v_i(S_i) = \sum_{j=0}^k f(j)|S_{ij}| = \sum_{j=0}^k (j+1)F = \Omega(k^2 F)$. Thus if we allocate each set S_i to buyer B_i for $i = 1, 2 \dots m$, the social welfare obtained is $\Omega(mk^2F)$, and hence OPT is $\Omega(mk^2F)$.

Consider now the arrival of some buyer B_i at time t. By our construction of the valuation function, B_i will either buy only shared items or buy only private items, but not both. If the buyer B_i were to buy shared items, and the price of each item is set at $f(j) \ge p > f(j+1)$, then B_i would pick up all remaining items in

$$\left(\bigcup_{0\le t\le j} S_{it}\right) \bigcup \left(\bigcup_{1\le \ell\ne i\le m}\ \bigcup_{0\le t\le (j-1)} S_{\ell t}\right)$$

Since $\sum_{0\le t\le j} Y^t \le 2Y^j$, the total price that B_i would pay to the seller is bounded by

$$\begin{aligned} f(j)(2Y^j + 2mY^{j-1}) &= 2(j+1)F + \frac{2m(j+1)F}{Y} \\ &= 2\left(1 + \frac{m}{Y}\right)(j+1)F \end{aligned}$$

We now consider the maximum revenue generated if B_i were to buy a subset of its private items. Note that when B_i arrives, all private items of B_i are still unsold. Suppose B_i were to buy private items. What is the maximum revenue we can get? For this, note that if the price of each item is $(j+1)F/jY^j$, then the utility from buying S'_{ij} is $(j+1)F - (j+1)F/j = (j+1)(j-1)F/j = (j-1/j)F$, and the utility from buying $S'_{i(j-1)}$ is $jF - (j+1)F/jY > (j-1/j^2)F > (j-1/j)F$, since $Y > 2j$.

So at this price, the set $S'_{i(j-1)}$ is preferred by B_i than S'_{ij}, and since the items in sets S'_{it} for $t > j$ have less value than the price, they are not even considered. For a greater price, the utility of $S'_{i(j-1)}$ must continue to dominate that of S'_{ij}, since the former has fewer items. So at most Y^{j-1} items are bought when the price is at least $(j+1)F/jY^j$, for all $j \geq 1$. This implies that the revenue obtained from B_i when she buys from her private items is at most $Y^j((j+1)F/jY^j) < 2F$.

Consider any ordering of buyers. If the price is ever set at more than $f(0)$, then no item is sold in that round, while if the price set is $f(k+1)$ or lower, all items are sold in that round and the revenue generated is at most $2mY^k f(k+1) = 2m(k+2)/Y$. Consider the first time when the price set in a round is at most $f(j)$ but greater than $f(j-1)$, for some $0 \leq j \leq k$. We call this round a *j-good sale*, and let B_i be the buyer. In a j-good sale, B_i may buy all remaining items in S_{it} for all $0 \leq t \leq j$, plus all items in S_{lt} for all $1 \leq \ell \leq m, \ell \neq i$ and $0 \leq t \leq j-1$, thus giving a revenue of at most $2(1+m/Y)(j+1)F \leq O((m/Y)kF)$

However, consider any time when a price in the range $(f(j-1), f(j)]$ appears again, and let B_l, $\ell \neq i$ be the buyer who faces this price. If B_l were to buy shared items, the only items that are valued higher than the price and still remaining are those in $S_{\ell j}$, since B_i took away whatever was remaining of $S_{\ell t}$ for all $t < j$. Note that the only reason the shared items could have given B_i a better utility was that the shared items had additive valuation, while the private items had XOS valuation, so she got no benefit in picking up multiple sets of private items. However, since only one feasible set $S_{\ell j}$ of the shared items is left, this advantage has vanished, and the revenue from B_ℓ is the same as the revenue if there were no shared items at all. As discussed above, the revenue from B_ℓ in this case is at most $2F$.

Finally, since a j-good sale can happen at most once for any $1 \leq j \leq (k+1)$, the total revenue generated fro all j-good sales is $O\left((mk^2F)/Y\right)$. The remaining rounds each give a revenue of at most $2F$, contributing in total $O(mF)$ to the revenue. Thus the revenue obtained by any dynamic uniform strategy, *for any ordering of buyers*, is $O\left((1+k^2/Y)mF\right)$.

Now since the maximum social welfare is $\Omega(k^2mF)$, the approximation factor achieved is bounded from below by $\Omega\left((k^2Y)/(k^2+Y)\right)$. For any $k > 10$, if we set $Y = k^2$ and $m = 2Y$, then $n = \Theta(Y^{k+1}) = k^{\Theta(k)}$, and the approximation factor is $\Omega(k^2)$. As $k = \Theta(\log n/\log\log n)$, we get that the smallest approximation factor that can be achieved is $\Omega\left((\log n/\log\log n)^2\right)$. ■

5. Dynamic Monotone Uniform Pricing

We now present a simple strategy that uses a monotonically decreasing uniform pricing for the items. When the number of buyers m is at least $2\log n$, the strategy gives an $O(\log^2 n)$-approximation to the revenue provided the buyers arrive in a uniformly random order, that is, all permutations of the buyers are equally likely to be the arrival order. As a corollary of this result, we conclude that if the buyers are identical, no matter the order in which they arrive, this pricing scheme gives an $O(\log^2 n)$-approximation. The strategy assumes that the seller knows the number of buyers m (and also OPT), and is *deterministic*. Knowing estimates of m and OPT up to constant factors are also sufficient for the performance of our strategy.

Let $k = \lfloor \log n \rfloor + 1$, and let $\gamma = 2^{k/m} \geq 1$. Thus $\gamma^m > 2n$. The strategy gives a good guarantee only when $m \geq \log n + 1$. The strategy is as follows: When the t^{th} buyer arrives, the seller prices all unsold items uniformly at

$$p[t] = \frac{\text{OPT}}{2\gamma^t} \quad .$$

Thus the price decreases with time. For $m = \omega(\log n)$, the relative decrease in the price for consecutive buyers is

$$\frac{p[t]-p[t+1]}{p[t]} = 1 - \frac{1}{\gamma} = \Theta\left(\frac{\log n}{m}\right) \quad ,$$

which tends to zero, and so the price decreases smoothly with time.

Theorem 9: Suppose $m \geq \log n + 1$, and suppose that the buyer valuations are subadditive. If the ordering of buyers in which they arrive is uniformly random (that is, all permutations are equally likely), then the expected revenue of the dynamic monotone uniform pricing scheme above is $\Omega(\frac{\text{OPT}}{\log^2 n})$.

Proof: Let $(T_1, T_2, \ldots T_m)$ be an optimal allocation of items to buyers $B_1, B_2 \ldots B_m$, who has valuation functions $v_1, v_2 \ldots v_m$ respectively, such that $\sum_{i=1}^m v_i(T_i) = \text{OPT}$ is the maximum social welfare. Also, let T_i^j be the subset of T_i that would be bought by B_i if it were shown only the items in T_i, and all items were uniformly priced at $\text{OPT}/\gamma^j = 2p[j]$.

Fix a buyer B_i. Let R_i be a random variable that denotes the revenue obtained by the seller from B_i. Let R'_i be a random variable that denotes the revenue obtained by the seller by selling items in T_i. Then, if R is a random variable that denotes the total revenue obtained by our strategy, we have $R = \sum_{i=1}^m R_i$ and $R \geq \sum_{i=1}^m R'_i$, so $R \geq \sum_{i=1}^m \frac{R_i + R'_i}{2}$.

Fix a permutation π of all buyers except B_i. We shall say that the event π occurs if these buyers arrive in the relative order given by π, with B_i arriving somewhere in between. We shall now compute $\mathbf{E}[R_i + R'_i|\pi]$.

Let π_j denote the permutation of all the buyers formed by inserting B_i after the $(j-1)^{th}$ but before the j^{th} position in π, whichever exists, for $1 \leq j \leq m$. That is B_i comes in as the j^{th} buyer in π_j. Let Z_i^j denote the number of items that were sold before the arrival of B_i when the arrival sequence of buyers is π_j. Note that Z_i^j is no longer a random variable once π_j is fixed, and neither are R_i and R'_i. Also note that

$\mathbf{Pr}[\pi_j|\pi] = 1/m$. Thus,

$$\mathbf{E}[R_i'|\pi] \geq \frac{1}{m}\sum_{j=1}^{m} p[j-1]|Z_i^j| \geq \frac{1}{m}\sum_{j=1}^{m} p[j]|Z_i^j| \ . \quad (3)$$

Let S_i^j be the set of items bought by B_i when the permutation of buyers is π_j, let $\mathcal{U}_i^j$ be the utility derived by B_i in this process, and let R_i^j be the revenue obtained from B_i. For $1 \leq j < m$, note that when the permutation is π_j, then when B_i arrives, the set S_i^{j+1} is also available, and B_i prefers S_i^j over this set at price $p[j]$. Thus

$$\begin{aligned}\mathcal{U}_i^j &= v_i(S_i^j) - p[j]|S_i^j| \geq v_i(S_i^{j+1}) - p[j]|S_i^{j+1}| \\ &= v_i(S_i^{j+1}) - p[j+1]|S_i^{j+1}| - (\gamma-1)p[j+1]|S_i^{j+1}| \\ &= \mathcal{U}_i^{j+1} - (\gamma-1)R_i^{j+1} \ .\end{aligned}$$

This implies that $R_i^{j+1} \geq \frac{1}{\gamma-1}(\mathcal{U}_i^{j+1} - \mathcal{U}_i^j)$, for $1 \leq j < m$. Also note that $v_i(S_i^1) - p[0]|S_i^1| \leq 0$, since no bundle of items can have value greater than $p[0] = \mathrm{OPT}$. So $v_i(S_i^1) - p[1]|S_i^1| - (\gamma-1)p[1]|S_i^1| = \mathcal{U}_i^1 - (\gamma-1)R_i^1 \leq 0$, or $R_i^1 \geq \frac{1}{\gamma-1}\mathcal{U}_i^1$. Thus, adding the terms R_i^t, we find that the terms telescope, and $\sum_{t=1}^{j} R_i^t \geq \frac{1}{\gamma-1}\left(\sum_{t=2}^{j}\left(\mathcal{U}_i^j - \mathcal{U}_i^{j-1}\right) + \mathcal{U}_i^1\right) = \mathcal{U}_i^j/(\gamma-1)$.

By Lemma 1, we have $v_i(T_i^j \setminus Z_i^j) \geq 2p[j]|T_i^j \setminus Z_i^j|$. So the utility of $T_i^j \setminus Z_i^j$ to buyer B_i at price $p[j]$ is at least $(2p[j]-p[j])|T_i^j \setminus Z_i^j| = p[j]|T_i^j \setminus Z_i^j|$, which is at most $\mathcal{U}_i^j$. Thus $\sum_{t=1}^{j} R_i^t \geq \mathcal{U}_i^j/(\gamma-1) = p[j]|T_i^j \setminus Z_i^j|/(\gamma-1)$. So we get

$$\begin{aligned}& m\sum_{j=1}^{m} R_i^j \geq \sum_{j=1}^{m}\sum_{t=1}^{j} R_i^t = \frac{1}{\gamma-1}\sum_{j=1}^{m} p[j]|T_i^j \setminus Z_i^j| \\ \Rightarrow & \sum_{j=1}^{m} R_i^j \geq \sum_{j=1}^{m}\frac{p[j]|T_i^j \setminus Z_i^j|}{(\gamma-1)m} \geq \Omega\left(\frac{p[j]|T_i^j \setminus Z_i^j|}{\log n}\right) \ .\end{aligned}$$

The last inequality uses the fact that $\gamma - 1 = \Theta\left(\frac{\log n}{m}\right)$.

Note that $\mathbf{E}[R_i|\pi] = \frac{1}{m}\sum_{j=1}^{m} R_i^j$. Combining with equation (3), we get that

$$\begin{aligned}\mathbf{E}[R_i + R_i'|\pi] &= \frac{1}{m}\Omega\left(\sum_{j=1}^{m} p[j]\left(\frac{|T_i^j \setminus Z_i^j|}{\log n} + |Z_i^j|\right)\right) \\ &\geq \Omega\left(\sum_{j=1}^{m}\frac{p[j]|T_i^j|}{m\log n}\right) \\ &= \Omega\left(\sum_{j=1}^{m}\frac{\mathrm{OPT}|T_i^j|}{(m\log n)\gamma^j}\right) \ .\end{aligned}$$

Using Lemma 3, we get that

$$\begin{aligned}\sum_{j=1}^{m}\frac{\mathrm{OPT}}{\gamma^j}|T_i^j| &\geq \frac{1}{\gamma-1}\left(v_i(T_i) - \frac{|T_i|\mathrm{OPT}}{\gamma^m}\right) \\ &\geq \frac{1}{\gamma-1}\left(v_i(T_i) - \frac{|T_i|\mathrm{OPT}}{2n}\right) \ .\end{aligned}$$

Again using the fact that $\gamma - 1 = \Theta\left(\frac{\log n}{m}\right)$, we get that

$$\mathbf{E}[R_i + R_i'|\pi] \geq \Omega\left(\frac{1}{\log^2 n}\left(v_i(T_i) - \frac{|T_i|\mathrm{OPT}}{2n}\right)\right) \ .$$

Since the right-hand-side of the above equation is independent of π, we get that the expected revenue is

$$\begin{aligned}\mathbf{E}[R] &= \sum_{i=1}^{m}\frac{\mathbf{E}[R_i + R_i']}{2} \geq \sum_{i=1}^{m}\Omega\left(\frac{v_i(T_i)}{\log^2 n} - \frac{|T_i|\mathrm{OPT}}{2n\log^2 n}\right) \\ &= \Omega\left(\frac{1}{\log^2 n}\left(\mathrm{OPT} - \frac{\mathrm{OPT}}{2}\right)\right) = \Omega\left(\frac{\mathrm{OPT}}{\log^2 n}\right) \ .\end{aligned}$$

■

6. Static Non-Uniform Pricing

Another approach to get around the weak performance barrier for static uniform pricing, is to consider static *non-uniform* pricing, which allows the seller to post different prices for different items but the prices remain unchanged over time. In Section 3 we showed that there exist instances with identical buyers where no static uniform pricing can achieve better than $2^{\Omega(\sqrt{\log n})}$-approximation even in the full information setting. Surprisingly, this hardness result breaks down if we consider non-uniform pricing using only two distinct prices.

6.1. Full Information Setting

We first introduce the (p,∞)-strategies, i.e. the seller posts price p for a subset of the items and posts ∞ for all other items. The intuition is by using this strategy the seller can prevent the buyers from buying certain items (high utility but low revenue) and thus achieve better revenue. The proof of the theorem below depends on the performance of the following dynamic monotone strategy. Let $k = \lceil \log n \rceil + 1$ and $m' = \lfloor m/k \rfloor$. Recall that $p_i = \mathrm{OPT}/2^i$ for $i = 1, 2, \cdots, k$. The seller posts a single price p_1 for the first m' buyers, then she posts a single price p_2 for the next m' buyers, and so on and so forth. We call each time period that the seller posts a fixed price a *phase*, and we call this strategy the k-phase monotone uniform strategy. The proof of Theorem 9 can be easily modified to show that this strategy gives $O(\log^2 n)$-approximation as well.

Theorem 10: In the full information setting, if $m \geq \log n + 1$, and all buyers share the same subadditive valuation function, then there exists a (p,∞)-strategy which obtains revenue at least $\Omega(\mathrm{OPT}/\log^3 n)$.

Proof: Given that the k-phase dynamic monotone uniform strategy for identical buyers obtains revenue at least $\Omega(\mathrm{OPT}/\log^2 n)$, at least one of the $k = \lceil \log n \rceil + 1$ phases contributes $1/k$ fraction of the revenue. Without loss of generality, assume the i^{th} phase contributes at least $\Omega(\mathrm{OPT}/k\log^2 n) = \Omega(\mathrm{OPT}/\log^3 n)$ revenue. Recall that $m' = \lfloor m/k \rfloor$. Suppose T is the set of items unsold at the beginning of phase i.

Consider the following (p, ∞)-strategy. The seller posts price $p = p_{i+1}$ for each item in T, and posts ∞ for all other items. Then when the first m' buyers come, they will behave the same as the m' buyers in phase i in the dynamic strategy scenario. So the revenue collected is at least $\Omega(\mathrm{OPT}/\log^3 n)$. ■

6.2. *Buyers with ℓ-XOS Valuations*

The above theorem shows a clear gap between the power of uniform pricing and the power of non-uniform pricing in the full information setting. However, it crucially uses the knowledge of the valuation function and the fact that all buyers are identical; information that is usually not known to the seller. Hence strategies in the limited information setting are more desirable in practice.

Fortunately, we find that considering static non-uniform pricing is also beneficial in the limited information setting. We first note that if the buyer order is randomized, then it is quite easy to get an $O\left(m \log n \log \mathrm{OPT} (\log\log \mathrm{OPT})^2\right)$ approximation using static uniform pricing, even with general valuations, and without the assumption of knowing OPT. This can be done as follows: Just focus on selling items to the first buyer. If B_i is the first buyer, and the algorithm knew the value $v_i(T_i)$, then using the single buyer (unlimited supply setting) algorithm in [4], the algorithm gets $\Omega(v_i(T_i)/\log n)$ in expectation from the first buyer, and we do not care what it gets from the other buyers. Thus the expected revenue of the algorithm is

$$\frac{1}{m}\left(\frac{\sum_{i=1}^{m} v_i(T_i)}{\log n}\right) = \frac{\mathrm{OPT}}{m \log n} \ .$$

This algorithm would have to guess $v_i(T_i) \le \mathrm{OPT}$ of the first buyer B_i, up to a constant factor, and can do so by incurring an additional factor of $O(\log \mathrm{OPT} (\log\log \mathrm{OPT})^2)$ as described in Lemma 4.

However, if we require a strategy to give guarantees on expected revenue against any order of buyers, and in particular an adversarial ordering, then static uniform pricing cannot give a better bound than $2^{\Omega(\sqrt{\log n})}$ even when there are only two buyers, with 3-XOS valuations. This is evident from the proof of Theorem 5. We now show a static non-uniform strategy which achieves polylogarithmic approximation if we assume the valuation functions are ℓ-XOS functions where ℓ is quasi-polynomial in n and the number of buyers is polylogarithmic, for all ordering of buyers.

Let $k = \lceil 2 \log n \rceil$. With probability half, the seller assigns a single price p randomly drawn from $\{p_1, p_2, \cdots, p_k\}$ to all items. With probability half, the seller assigns one of $p_1, p_2, \cdots, p_{k+1}$ uniformly at random for each item. Here $p_i = \mathrm{OPT}/2^i$. The price assignment remains unchanged over time.

Theorem 11: For m buyers with ℓ-XOS valuations functions, the expected revenue of the above strategy is

$$\Omega\left(\frac{\mathrm{OPT}}{m \log \ell \log^3 n}\right) \ .$$

The proof of the above theorem is deferred to the full version of the paper.

References

[1] G. Aggarwal, T. Feder, R. Motwani, and A. Zhu, "Algorithms for multi-product pricing," in *ICALP*, 2004, pp. 72–83.

[2] M.-F. Balcan and A. Blum, "Approximation algorithms and online mechanisms for item pricing," in *ACM Conference on Electronic Commerce*, 2006, pp. 29–35.

[3] M.-F. Balcan, A. Blum, J. D. Hartline, and Y. Mansour, "Mechanism design via machine learning," in *FOCS*, 2005, pp. 605–614.

[4] M.-F. Balcan, A. Blum, and Y. Mansour, "Item pricing for revenue maximization," in *ACM Conference on Electronic Commerce*, 2008, pp. 50–59.

[5] P. Briest, M. Hoefer, and P. Krysta, "Stackelberg network pricing games," in *STACS*, 2008, pp. 133–142.

[6] P. Briest and P. Krysta, "Single-minded unlimited supply pricing on sparse instances," in *SODA*, 2006, pp. 1093–1102.

[7] P. Cramton, Y. Shoham, and R. Steinberg, *Combinatorial Auctions*. MIT Press, 2005.

[8] E. D. Demaine, M. T. Hajiaghayi, U. Feige, and M. R. Salavatipour, "Combination can be hard: approximability of the unique coverage problem," in *SODA*, 2006, pp. 162–171.

[9] S. Dobzinski, "Two randomized mechanisms for combinatorial auctions," in *APPROX-RANDOM*, 2007, pp. 89–103.

[10] S. Dobzinski, N. Nisan, and M. Schapira, "Truthful randomized mechanisms for combinatorial auctions," in *STOC*, 2006, pp. 644–652.

[11] U. Feige, "On maximizing welfare when utility functions are subadditive," in *STOC*, 2006, pp. 41–50.

[12] A. Fiat, A. V. Goldberg, J. D. Hartline, and A. R. Karlin, "Competitive generalized auctions," in *STOC*, 2002, pp. 72–81.

[13] A. V. Goldberg, J. D. Hartline, and A. Wright, "Competitive auctions and digital goods," in *SODA*, 2001, pp. 735–744.

[14] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, "On profit-maximizing envy-free pricing," in *SODA*, 2005, pp. 1164–1173.

[15] J. Hartline and A. Karlin, "Profit maximization in mechanism design," in *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Eds. Cambridge University, 2007, ch. 13.

[16] R. Lavi and C. Swamy, "Truthful and near-optimal mechanism design via linear programming," in *FOCS*, 2005, pp. 595–604.

[17] D. J. Lehmann, L. O'Callaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *J. ACM*, vol. 49, no. 5, pp. 577–602, 2002.

[18] R. B. Myerson, "Optimal auction design," *Mathematics of Operations Research*, vol. 6, pp. 58–73, 1981.