

Item Pricing for Revenue Maximization

Maria-Florina Balcan*

Avrim Blum†

Yishay Mansour‡

ABSTRACT

We consider the problem of pricing n items to maximize revenue when faced with a series of unknown buyers with complex preferences, and show that a simple pricing scheme achieves surprisingly strong guarantees.

We show that in the unlimited supply setting, a random single price achieves expected revenue within a logarithmic factor of the total social welfare for customers with *general* valuation functions, which may not even necessarily be monotone. This generalizes work of Guruswami et. al [18], who show a logarithmic factor for only the special cases of single-minded and unit-demand customers.

In the limited supply setting, we show that for *subadditive* valuations, a random single price achieves revenue within a factor of $2^{O(\sqrt{\log n \log \log n})}$ of the total social welfare, i.e., the optimal revenue the seller could hope to extract even if the seller could price each bundle differently for every buyer. This is the best approximation known for any item pricing scheme for subadditive (or even submodular) valuations, even using multiple prices. We complement this result with a lower bound showing a sequence of subadditive (in fact, XOS) buyers for which any single price has approximation ratio $2^{\Omega(\log^{1/4} n)}$, thus showing that single price schemes cannot achieve a polylogarithmic ratio. This lower bound demonstrates a clear distinction between revenue maximization and social welfare maximization in this setting, for which [12, 10] show that a fixed price achieves a logarithmic approximation in the case of XOS [12], and more generally subadditive [10], customers.

*School of Computer Science, Carnegie Mellon University. Supported in part by NSF grant CCF-0514922, by an IBM Graduate Fellowship, and by a Google Research Grant.

†School of Computer Science, Carnegie Mellon University. Supported in part by NSF grant CCF-0514922 and by a Google Research Grant

‡Google Inc. and School of Computer Science, Tel Aviv University. This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a grant no. 1079/04 from the Israel Science Foundation, by a grant from United States-Israel Binational Science Foundation (BSF), and an IBM faculty award. This publication reflects the authors' views only.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'08, July 8–12, 2008, Chicago, Illinois, USA.

Copyright 2008 ACM 978-1-60558-169-9/08/07 ...\$5.00.

We also consider the multi-unit case examined by [11] in the context of social welfare, and show that so long as no buyer requires more than a $1 - \epsilon$ fraction of the items, a random single price now does in fact achieve revenue within an $O(\log n)$ factor of the maximum social welfare.

Categories and Subject Descriptors

F.2 [Analysis of Algorithms and Problem Complexity]: General

General Terms

Algorithms, Theory, Economics.

Keywords

Approximation Algorithms, Combinatorial Auctions

1. INTRODUCTION

Item pricing is one of the most fundamental problems in Economics, and to a large extent describes today's trading practices. From a theoretical perspective, the issue of market equilibrium prices has received enormous attention over the years. In this work we do not consider the equilibrium issues of market pricing, but rather concentrate on the basic problem of revenue maximization. We consider a single seller whose goal is to maximize his revenue, who must set prices on items before the arrival of a sequence of customers with complex, unknown preferences (e.g., think of a store or a yard sale). We prove that a simple posted single pricing scheme produces the best revenue guarantees known for this problem for two important classes of settings: buyers with *general* valuations for the case of items in unlimited supply, and buyers with *subadditive* valuations for the case of items in limited supply. Note that no good approximation is possible for general valuations with limited supply.

Formally, the problem we analyze is the following. We consider a single seller who has n items each in limited or unlimited supply. We assume there are m buyers with quasi-linear utilities who arrive in an arbitrary order and who have unknown and potentially highly complex valuations over subsets of these items.¹ We consider the revenue maximization objective—the goal of the seller is to make as much money as possible—and analyze a natural and ubiquitous setting in which the seller first posts prices on the items and then buyers enter one at a time and purchase whatever subset of the remaining items they want most. Since prices are fixed before buyers

¹Quasi-linear utilities means that buyers prefer the set maximizing the difference between its cost and its value. We assume the usual oblivious adversary model in which the sequence and valuations of buyers is determined in advance of any randomization made by the seller.

arrive, all upper bounds we present for our pricing scheme will also apply to the problem of designing *truthful mechanisms* with revenue guarantees in the setting of combinatorial auctions. Note that while much work on combinatorial auctions considers bundle-pricing mechanisms (such as based on VCG), the vast majority of transactions in today’s world are conducted via pricing on items, and thus it is important to understand what guarantees are possible in such a setting. As an upper bound on the revenue that the seller can hope to extract from the buyers we use the optimum social welfare, which is the maximum possible sum of buyers’ valuations in any allocation. This is the most revenue the seller could extract even if the seller could price each bundle differently for every buyer.

In the limited supply setting, we consider important classes of valuations functions which have previously been studied in the context of maximizing social welfare: submodular, XOS, and more generally, subadditive valuation functions [14, 24, 22, 12, 26, 10]. We show here that for buyers with subadditive valuation functions, a random single price achieves revenue within a $2^{O(\sqrt{\log n \log \log n})}$ factor of the maximum social welfare. We complement this result with a lower bound showing a sequence of subadditive (in fact, XOS) buyers for which any single price has approximation ratio $2^{\Omega(\log^{1/4} n)}$, thus showing that single prices cannot achieve a poly-logarithmic ratio. Moreover, this lower bound holds even if the price is determined based on advance knowledge of the order and valuations of the buyers. The construction in this lower bound demonstrates a clear distinction in this setting between revenue maximization and social welfare maximization, for which [12, 10] show that a fixed price achieves a logarithmic approximation in the case of XOS [12], and more generally subadditive [10], customers. We also show that even if we assume buyers arrive in a *random* order, there exists a set of buyers for which a $2^{\Omega(\log^{1/4} n)}$ lower bound still holds. Note that our $2^{O(\sqrt{\log n \log \log n})}$ upper bound is the best approximation known for *any* item pricing scheme for sub-additive buyers, even if assigning different prices to different items is allowed. We also show that for a special case we call *simple sub-modular valuations* (which generalizes unit-demand, additive, and submodular symmetric valuations [23]), a random single price does in fact achieve revenue within a logarithmic factor of the optimum social welfare.

In the unlimited supply setting, we show that for buyers with general valuation functions, there exists a single price one can assign to all the items such that the retailer achieves revenue within a logarithmic factor of the total social welfare (and furthermore this holds in expectation if that price is chosen at random from an appropriate distribution).² Our main result in the unlimited supply setting (Theorem 9) turns out to provide a useful and convenient structural characterization needed for proving the desired approximation for subadditive valuations in the limited supply case.

Finally, we consider the multi-unit auctions setting [11, 23] where we have only one item and multiple copies of it, but buyers have *arbitrarily complicated* valuation functions over the number of copies received. We show that under the assumption that the optimal allocation gives at most a $(1 - \epsilon)$ fraction of the items to any one buyer, our single pricing scheme achieves a logarithmic approximation in this setting as well.

Related work: Guruswami et al. [18] show that in an unlimited supply combinatorial auction, if customers are unit-demand or single-minded,³ then a random single price achieves expected

²An early version of our paper, with just this result, appears as CMU Tech Report CMU-CS-07-111 [4].

³A single-minded buyer is one who places some value v on a single set S or any superset of S , and value 0 on any set that does not

contain S . A unit-demand buyer is one who has separate values v_j on each item j , and values any given set S at $\max_{j \in S} v_j$.

revenue within a logarithmic factor of the total social welfare. In this paper we show the restriction to single-minded or unit-demand valuations is not required: in particular, we show that a random single price achieves this guarantee for buyers with *general* valuation functions over bundles, which may not even necessarily be monotone. Moreover, no item pricing scheme (even one that assigns each item a different price) can do better, even for a single customer whose valuations are known to the seller.

Our result for the unlimited supply case appears in an earlier version of this paper [4]. It was also discovered independently in a different context by Briest et al. [6] who study single price schemes in a network setting. In their setting, a buyer has certain subgraphs of the network it is interested in purchasing. A seller, who owns the network, first prices the edges and then the buyer purchases the cheapest subgraph it is interested in. They show that a single fixed price for all the edges guarantees the seller a revenue within logarithmic factor of the highest possible revenue.

As mentioned above, the setting we analyze can be viewed as an online version of combinatorial auctions. The literature on combinatorial auctions is extensive and spans the fields of Economics, Operations Research and Computer Science [8, 26]. In the Computer Science community, there have recently been two main threads of work: designing good algorithms [18, 1, 7, 2, 9, 20] and mechanisms [15, 16, 3, 19] for *revenue maximization* in the unlimited supply setting,⁴ and designing algorithms [14, 23] and computationally efficient mechanisms [24, 22, 12, 26, 10] to optimize *social welfare* in the limited supply setting. For a detailed related work section see Appendix B, and for excellent recent overviews see [5, 26, 21, 19].

In the context of social welfare maximization, work most related to ours is that of Dobzinski et al. [12, 10], who show that in the limited supply setting, a fixed price achieves a logarithmic approximation in the case of XOS or subadditive [12, 10] customers. In our work we analyze its power for maximizing revenue.

Structure of this paper: This paper is organized as follows. We start with terminology and formal definitions in Section 2. We then present our results for the limited-supply setting in Section 3: we show that for subadditive buyers, a random single price (with buyers arriving in an arbitrary order) achieves a ratio $2^{O(\sqrt{\log n \log \log n})}$, along with a $2^{\Omega(\log^{1/4} n)}$ lower bound for any single-price scheme (even if buyers’ valuations are known in advance and buyers arrive in a random order). We then analyze the unlimited supply setting in Section 4 and prove that a random single price achieves a logarithmic ratio for buyers with *general* valuation functions. We consider multi-unit auctions as well as a special case of submodular valuations in Section 5 and finish with a discussion in Section 6.

2. PRELIMINARIES

We consider the following setting. A single seller has a set J of n items, and there is a sequence B of m buyers or customers, who are interested in buying the items.

Each buyer $i \in B$ has a private valuation $\mathbf{v}_i(S)$ for each bundle $S \subseteq J$ of items, which measures how much receiving bundle S would be worth to him. The utility of the buyer $i \in B$ for purchasing the set T is $\mathbf{u}_i(T) = \mathbf{v}_i(T) - \sum_{j \in T} p_j$, where p_j is the price of item $j \in J$. (For a single price p we have $\mathbf{u}_i(T) = \mathbf{v}_i(T) - |T|p$.) That is, we assume that a buyer’s utility is quasi-linear. Note that we implicitly assume that there are no externalities since the buyers’ utilities are completely determined by the

contain S . A unit-demand buyer is one who has separate values v_j on each item j , and values any given set S at $\max_{j \in S} v_j$.

⁴Note that the social welfare objective is trivial in this setting: one simply gives everything away for free.

set of items purchased and the price paid; buyers do not care about the happiness of the *other* buyers, for instance. Finally, given the valuation functions of the buyers, we define $H_i = \max_S (v_i(S))$ and $H = \max_{i,S} (v_i(S))$.

A buyer's valuation function might be quite complex since there are 2^n possible bundles, but we make the minimal (standard) assumption that given a vector \mathbf{p} of item prices and a set of items $J' \subset J$, the customer can determine the subset $T \subset J'$ it most wants at those prices. Formally, the buyer has a demand oracle [26], such that $DemandPrices(i, \mathbf{p}, J')$ returns the set

$$T = \arg \max_{S \subset J'} \mathbf{u}_i(S).$$

We analyze important classes of valuation functions which have received substantial attention in the combinatorial auctions literature [14, 10, 12, 5]: submodular, XOS, and more generally, subadditive valuation functions. A valuation function \mathbf{v} is submodular if $\mathbf{v}(S \cup T) + \mathbf{v}(S \cap T) \leq \mathbf{v}(S) + \mathbf{v}(T)$, for all $S, T \subseteq J$. A valuation \mathbf{v} is subadditive if $\mathbf{v}(S \cup T) \leq \mathbf{v}(S) + \mathbf{v}(T)$, for all $S, T \subseteq J$. Between these two classes (submodular and subadditive) lies the class of "XOS" valuations. A valuation \mathbf{v} is XOS if there are additive valuations $\{a_1, \dots, a_k\}$ such that $\mathbf{v}(S) = \max_k a_k(S)$, for all $S \subseteq J$. Submodular is strictly more restrictive than XOS which is strictly more restrictive than subadditive.

We study both the limited supply setting, where without loss of generality we may assume that exactly one copy of each item is available, and the unlimited setting where the number of copies of each item is as large as the number of buyers. It is assumed that each buyer wants at most one copy of each item.

The pricing scheme we analyze throughout most of this paper is a *single posted price* mechanism. The seller starts by choosing a single posted price p at random from an appropriate distribution. The buyers then arrive in an arbitrary order, and each buyer at his turn buys his most preferred set at the given price of p per item. (The order of buyers does not matter in the unlimited supply case.) The revenue of the seller is then the total number of items purchased times p . Clearly, since the posted price p is chosen independently of the buyers valuations, this scheme falls into the category of incentive-compatible mechanisms. In fact, there is no communication from the buyers, other than selecting the subset of items they will purchase. (Note that even if we allow buyers to purchase at multiple different times, since the price does not change, his best policy is to buy the set of items that he desires when he first arrives.)

In the limited supply setting, where one copy of each item is available, we say that an allocation T_1, \dots, T_n is feasible if $T_i \cap T_j = \emptyset$ for $i \neq j$. The *social welfare* of a feasible allocation T_1, \dots, T_n is $\sum_i v_i(T_i)$. The *social optimum* is the value of the allocation which maximizes the social welfare, which upper bounds the seller's revenue under any mechanism (even if the seller can have a different price for each bundle and buyer). In the unlimited supply setting the *social optimum* is $\sum_{i=1}^m H_i$. Our bounds will compare the expected revenue of the seller to the social optimum.

For simplicity, we assume throughout the paper that we know H , and we remark on how to overcome exactly knowing H in Appendix A.

3. LIMITED SUPPLY: CUSTOMERS WITH SUBADDITIVE VALUATIONS

In this section we consider the case of buyers with subadditive valuations and analyze the single posted price mechanism *RANDOM Single Price* (Algorithm 1).

We start by proving our main upper bound for revenue maximization in the limited-supply setting, showing that *RANDOM Sin-*

Algorithm 1 *RANDOM Single Price*

Input: $H = \max_{i,S} (v_i(S))$, and s a parameter.

Step 1 Let $q_l = \frac{H}{2^{l-1}}$, for $l \in \{1, \dots, s\}$.

Step 2 Pick a posted price p uniformly at random in $\{q_1, \dots, q_s\}$;

Step 3 Buyers arrive in an arbitrary order and purchase their most preferred bundle. I.e., set $R = J$ and do:

Procedure *Generate Allocation* (p)

For buyer i , let $S_i = DemandPrice(i, p, R)$.

Allocate S_i to buyer i and charge it $p|S_i|$

Let $R = R \setminus S_i$.

Step 4 The seller has a remainder set of items R .

gle Price achieves a $2^{O(\sqrt{\log n \log \log n})}$ approximation to the social optimum, assuming the buyer's valuations are subadditive. We begin with a definition from [12].

DEFINITION 1. An allocation $S = (S_1, \dots, S_m)$ is supported at price p if, for each buyer i and for every possible bundle $W_i \subseteq S_i$, it holds that $\mathbf{v}_i(W_i) \geq p|W_i|$.

Before presenting the proof of our main result we first give two useful lemmas. The first lemma states that if the valuation functions are subadditive, then for every possible allocation it is possible to find a "contained" allocation and a price p that supports it such that if buyers purchased the supported allocation it would produce revenue comparable to the welfare of the original allocation. (This is based on a result we prove in the next section.) Unfortunately buyers left to their own devices might not purchase the supported allocation, however. The second lemma, though, states that if *Generate Allocation* is run at price $p/2$, then the allocation produced at least will have large social welfare, even if the revenue is not so high.

LEMMA 1. Assume that \mathbf{v}_i are subadditive. Let T_1, \dots, T_m be an arbitrary feasible allocation and let $\alpha = \frac{1}{4 \log(2n^2)}$. There exists a price p and subsets $L_i \subseteq T_i$ such that L_1, \dots, L_m is an allocation supported at price p and furthermore

$$\sum_{i=1}^m \mathbf{v}_i(T_i) \geq \sum_{i=1}^m p|L_i| \geq \alpha \sum_{i=1}^m \mathbf{v}_i(T_i).$$

The proof follows from Theorem 9 (which will be presented in Section 4) and properties of subadditive valuation functions. In particular, we set $L_i = DemandPrice(i, p, T_i)$. The above lemma suggests that if the seller can present every buyer a different set T_i , then its revenue would be close to the social welfare of the T_i s. (For a full proof see Appendix A.)

The following lemma states that we have a reasonable chance that the produced allocation has a high social value. The proof follows along the lines of [12]. For completeness, we include it in Appendix A.

LEMMA 2. Let L_1, \dots, L_m be an allocation supported at price p . Let S_1, \dots, S_m be the allocation produced by *Generate Allocation* with the price parameter $p/2$. Then:

$$\sum_{i=1}^m \mathbf{v}_i(S_i) \geq \sum_{i=1}^m (p/2)|L_i|.$$

We now give our main upper bound for revenue maximization in limited-supply setting, showing that *RANDOM Single Price* achieves a $2^{O(\sqrt{\log n \log \log n})}$ approximation to the social optimum, assuming the buyers' valuations are subadditive. The high level idea of the proof is the following. We first show that if we could limit the buyers to purchase only subsets of their assigned bundle in the social welfare-maximizing allocation, then we do achieve a logarithmic approximation using a random price p . However, the problem is that the buyers, when left to their own devices (i.e., when they are allowed to purchase any subset they want), given a price p might buy a totally different subset, hurting the utility of subsequent buyers. Here we might have two outcomes. The easy case is when the number of items sold is sufficiently large, in which case we are done. (Since we have a single price, the seller does not care which items are sold, only how many.) The more difficult case is when the number of items sold is fairly small. In this case we show that we must have a small subset of the items such that its social welfare is not too much less than the original. We then use the same argument recursively on this smaller subset, and the recursive argument has to be complete before we get to an empty set.

THEOREM 3. *Assume that all the buyers have subadditive valuations, and let $s = \lfloor \log_2(2n^2) \rfloor$. The expected revenue of the *RANDOM Single Price mechanism* is $OPT/2^{O(\sqrt{\log n \log \log n})}$, where OPT is the social optimum.*

PROOF. Consider $\beta > 0$, and let $\tilde{\alpha} = \frac{\alpha}{2}$ where $\alpha = \frac{1}{4 \log(2n^2)}$. Let $T^1 = (T_1^1, \dots, T_m^1)$ be an allocation that maximizes the total social welfare. By Lemma 1 we know that there exists p_1 and an allocation $L^1 = (L_1^1, \dots, L_m^1)$, $L_i^1 \subseteq T_i^1$, such that L^1 is supported at price p_1 and $\sum_{i=1}^m p_1 |L_i^1| \geq \alpha \sum_{i=1}^m \mathbf{v}_i(T_i^1) = \alpha OPT$. Let S_1^1, \dots, S_m^1 be the allocation produced by *Generate Allocation* when run with the price parameter $p_1/2$. By Lemma 2 we know $\sum_{i=1}^m \mathbf{v}_i(S_i^1) \geq \sum_{i=1}^m (p_1/2) |L_i^1|$. So,

$$\sum_{i=1}^m \mathbf{v}_i(S_i^1) \geq (\alpha/2) \sum_{i=1}^m \mathbf{v}_i(T_i^1) = \tilde{\alpha} OPT.$$

If we additionally have $\sum_{i=1}^m |S_i^1| \geq \beta \sum_{i=1}^m |L_i^1|$, then the profit of our algorithm at price $p_1/2$ is at least $\tilde{\alpha}\beta OPT$. Otherwise, if $\sum_{i=1}^m |S_i^1| < \beta \sum_{i=1}^m |L_i^1|$, let us denote by $T^2 = (T_1^2, \dots, T_m^2)$ the allocation $S^1 = (S_1^1, \dots, S_m^1)$. We now repeat this process recursively on T^2 .

In general, assume inductively that at iteration l of the argument we have allocation T^l with $|T^l| \leq \beta^{l-1}n$ and $\sum_{i=1}^m \mathbf{v}_i(T_i^l) \geq \tilde{\alpha}^{l-1} OPT$. We know that there exists p_l and an allocation $L^l = (L_1^l, \dots, L_m^l)$ with $L_i^l \subseteq T_i^l$ such that L^l is supported at price p_l and $\sum_{i=1}^m p_l |L_i^l| \geq \alpha \sum_{i=1}^m \mathbf{v}_i(T_i^l) \geq \alpha \tilde{\alpha}^{l-1} OPT$. Let S_1^l, \dots, S_m^l be the allocation produced by *Generate Allocation* when run with the price parameter $p_l/2$. From Lemma 2 we know

$$\sum_{i=1}^m \mathbf{v}_i(S_i^l) \geq \sum_{i=1}^m (p_l/2) |L_i^l|.$$

So,

$$\sum_{i=1}^m \mathbf{v}_i(S_i^l) \geq (\alpha/2) \sum_{i=1}^m \mathbf{v}_i(T_i^l) = \tilde{\alpha}^l OPT.$$

If we additionally have $\sum_{i=1}^m |S_i^l| \geq \beta \sum_{i=1}^m |L_i^l|$, then the profit of our algorithm at price $p_l/2$ is at least $\tilde{\alpha}^l \beta OPT$. Otherwise, if $\sum_{i=1}^m |S_i^l| < \beta \sum_{i=1}^m |L_i^l|$, let us denote by $T^{l+1} = (T_1^{l+1}, \dots, T_m^{l+1})$ the allocation $S^l = (S_1^l, \dots, S_m^l)$, and we have $|T^{l+1}| \leq \beta^l n$, maintaining the induction.

Consider $\beta = \frac{1}{n^\epsilon}$. Then this process can continue for at most $l = 1/\epsilon$ rounds and thus our argument above implies that the profit of our algorithm is at least $OPT / (n^\epsilon (8 \log 2n^2)^{\frac{1}{\epsilon}+1})$. Setting $\epsilon = \sqrt{\frac{\log \log n}{\log n}}$, we obtain the desired competitive ratio of $2^{O(\sqrt{\log n \log \log n})}$. \square

Note that in the argument above, the prices p_0, \dots, p_l are monotonically increasing. To better understand the argument and to motivate the lower bound given in Theorem 4 below, consider the following interesting example with just one customer, whose valuation function is defined as follows. Partition the n items into sets $S_0, S_1, S_2, \dots, S_t$ where set S_i has $n_i = n_{i-1}/X$ items (so $n = n_0 + n_0/X + \dots + n_0/X^t$; $X \gg 1$ will be determined later). The valuation of our buyer is additive over items *within* any given set S_i and then the maximum over the sets S_i (so it is XOS). Assume that the value for each of the items in S_i is $v_i = ((X+1+\epsilon)/2)^i$, so the value of the i -th bundle $n_i v_i$ is approximately $n/2^i$. The set of highest value is S_0 , so suppose that as in the argument above, the seller chooses price $p = v_0/2$. Then, however, the buyer will purchase S_1 instead because $n_0(v_0 - p) < n_1(v_1 - p)$, and then the seller makes a factor X less revenue. On the other hand, if the seller chooses price $p = v_1/2$ then the same reasoning shows that the buyer will instead buy S_2 , and so on. This implies that if we limit ourselves to prices of the form $p_i = v_i/2$, as our argument does, the best single price is $v_t/2$ which produces a revenue of only approximately $n/2^t$. For $t = \sqrt{\log n}$ and $X = 2^t = 2^{\sqrt{\log n}}$ we get a loss of $\Omega(2^{\sqrt{\log n}})$. Note however that since $m = 1$, we know that there is a single price which is $\log n$ competitive — see Theorem 9; e.g. the price $v_1/4$ would provide the desired ratio in this example.

This example raises the question of whether an alternative analysis could yield a better upper bound. We prove below a surprising lower bound showing that even with just two buyers, due to the *interaction* between them, one cannot achieve a polylog(n) ratio by any single-price algorithm, even if the buyers' valuations are known in advance. This demonstrates a clear distinction between the goals of revenue and social welfare maximization in the limited supply, subadditive (or XOS) setting.

THEOREM 4. *There exists a set of buyers with XOS valuations, and an ordering of the buyers, such that any single posted price (even chosen based on the buyers' valuations) produces revenue at most $OPT/2^{(\log n)^{1/4}}$.*

PROOF. Let $X = 2^{(\log n)^{1/4}}$. Our goal is to show that no single price can beat the ratio X . As we will see, it suffices to consider only two buyers. The construction is inspired by the example above, though it is a bit more intricate.

Let us partition the items into sets $S_0, S_1, S_2, \dots, S_t$ where set S_i has n/X^i items and $t = \log_X(n) = (\log n)^{3/4}$ (so, technically, the total number of items is slightly larger than n). The valuation of buyer 2 will be additive over items *within* any given set S_i and then max over the bundles. Notice that this implies that at any given price, buyer 2 will purchase items from at most one set S_i .

We will define buyer 2 so that $v_2(S_i) = (1 - 1/\sqrt{\log n})v_2(S_{i-1})$. The high level idea of the construction is that buyer 2's valuations contain almost all of the total social welfare (the auctioneer cannot hope to make sufficient revenue from buyer 1). However, when prices are such that buyer 2 would ordinarily purchase S_i , and the price is high enough so that this constitutes substantial revenue, buyer 1 (who arrives first) purchases just enough of set S_i , $i < t$ to make buyer 2 choose S_{i+1} instead, reducing total revenue obtained from buyer 2 by a factor of X . In addition, the final set S_t has too low total valuation since $(1 - 1/\sqrt{\log n})^t \approx e^{-(\log n)^{1/4}}$. Thus

the auctioneer cannot possibly receive enough revenue by having buyer 2 purchase the final set.

In order to make this work, we need for buyer 1 to be able to cause buyer 2 to switch to S_{i+1} by purchasing only a $1/X$ fraction of S_i (if buyer 1 purchased more than this fraction of S_i , then the auctioneer would make too much revenue from buyer 1). We do this by defining buyer 2's valuations as follows. Let S'_i denote the first $|S_i|/X$ elements within S_i and let $L_i = X^i(1 - 1/\sqrt{\log n})^i$. We define buyer 2 to have value L_i on each of the items in $S_i - S'_i$ and value $L_i(X - 1)/(\sqrt{\log n} - 1)$ on each of the items in S'_i . So, for instance, buyer 2 has value 1 on the items in $S_0 - S'_0$ and value $(X - 1)/(\sqrt{\log n} - 1)$ on items in S'_0 . Thus,

$$v_2(S_i - S'_i) = n(1 - 1/X)(1 - 1/\sqrt{\log n})^i$$

and

$$v_2(S'_i) = n(1 - 1/X)(1 - 1/\sqrt{\log n})^{i-1}(1/\sqrt{\log n}).$$

Putting these together we get:

$$\begin{aligned} v_2(S_i) &= n(1 - 1/X)(1 - 1/\sqrt{\log n})^{i-1} \\ &= v_2(S_{i-1} - S'_{i-1}). \end{aligned}$$

In particular, the key points of this construction are the following: (a) $v_2(S_{i+1}) = v_2(S_i)(1 - 1/\sqrt{\log n})$, (b) $v_2(S_{i+1}) = v_2(S_i - S'_i)$, and (c) $|S'_i| = |S_{i+1}| = |S_i|/X$. In addition, buyer 2's valuations are such that at any given price, it prefers to purchase one of the sets S_i in its entirety.

We will define buyer 1's valuations to be nonzero only over the sets S'_0, S'_1, \dots, S'_i , and a max of sums just like buyer 2. The property we want from buyer 1 is that it should purchase S'_i when the price is in the range $[L_i/X, 3L_i/\sqrt{\log n}]$ and furthermore it should never produce much revenue for the seller. The reason we care only about this range is that below the lower end, we do not care if buyer 2 purchases S_i because the revenue to the auctioneer will be too low (less than $v_2(S_i)/X$). Above the upper end, we can show that buyer 2 prefers $S_{i+1} - S'_{i+1}$ to S_i so he will not purchase S_i even if buyer 1 purchases S'_{i+1} . Specifically, the price at which buyer 2 is indifferent between sets S_i and $S_{i+1} - S'_{i+1}$ is the price p such that:

$$\begin{aligned} n\left(1 - \frac{1}{X}\right)\left(1 - \frac{1}{\sqrt{\log n}}\right)^{i-1} - p\frac{n}{X^i} \\ = n\left(1 - \frac{1}{X}\right)\left(1 - \frac{1}{\sqrt{\log n}}\right)^{i+1} - p\frac{n}{X^{i+1}}\left(1 - \frac{1}{X}\right). \end{aligned}$$

Solving, we have

$$\begin{aligned} p &= X^i\left(1 - \frac{1}{X}\right)\left(1 - \frac{1}{\sqrt{\log n}}\right)^{i-1}\left(1 - \left(1 - \frac{1}{\sqrt{\log n}}\right)^2\right) \\ &\quad \left(\frac{X}{X - (1 - \frac{1}{X})}\right) \\ &\leq X^i\left(1 - \frac{1}{\sqrt{\log n}}\right)^{i-1}\left(\frac{2}{\sqrt{\log n}} - \frac{1}{\log n}\right) \\ &\leq L_i\left(\frac{1}{\sqrt{\log n}}\right)\left(\frac{2\sqrt{\log n} - 1}{\sqrt{\log n} - 1}\right), \end{aligned}$$

which is at most $3L_i/\sqrt{\log n}$ for $n \geq 16$. Moreover, we do not need to worry that buyer 2 might prefer just S'_i to $S_{i+1} - S'_{i+1}$ since both sets have approximately the same size and yet items in S'_i are much less valuable to buyer 2 than items in $S_{i+1} - S'_{i+1}$.

It remains to precisely define buyer 1. To get the desired behavior for this buyer we set its value on each of the items in S'_i to $4L_i$. So $v_1(S'_i)$ is $4(n/X)(1 - 1/\sqrt{\log n})^i$. We can now check that buyer 1 has the desired purchasing behavior: he purchases S'_i when prices are in the range $[L_i/X, 3L_i/\sqrt{\log n}]$, and yet he does not provide enough profit to the auctioneer. To see that this is true just notice that for all $p = \alpha L_i$, $\alpha \in [1/X, 3/\sqrt{\log n}]$ we have

$4(n/X)(1 - 1/\sqrt{\log n})^i - \alpha(n/X^{i+1})X^i(1 - 1/\sqrt{\log n})^i \geq 4(n/X)(1 - 1/\sqrt{\log n})^j - \alpha(n/X^{j+1})X^j(1 - 1/\sqrt{\log n})^j$. for all $j \neq i$. It's easy to see that this is true always for $j > i$ and it is also true when $j < i$ for large enough n . Thus, buyer 1 causes buyer 2 to purchase a set with a factor X less revenue to the auctioneer than it would have purchased without the presence of buyer 1, at any price for which the auctioneer would have made substantial revenue. This shows that any single price results in revenue that is a factor X worse than the total social welfare. \square

Note: It is easy to modify our example in Theorem 4 so that no buyer has a significant fraction of the social welfare. Specifically, we just need to make n "copies" of the example in Theorem 4, each on a completely disjoint set of items (so there are $2n$ buyers now, and $O(n^2)$ items total), such that each buyer has valuation 0 for all items not from their own set. Then clearly all buyers have valuation close to the average.

Clearly, the lower bound in Theorem 4 depends on a specific adversarial ordering of the buyers. However, even if we assume buyers arrive in a random order, there exists a set of buyers for which a $2^{\Omega(\log^{1/4} n)}$ bound still holds. Specifically:

THEOREM 5. *There exists a set of buyers with XOS valuations, such that any single posted price (even chosen based on the buyers' valuations) produces an expected revenue at most $(m/X + 1/m)OPT$, even under a random ordering of the buyers, where $X = 2^{(\log n)^{1/4}}$. Setting $m = \sqrt{X}$ we have a lower bound of $2^{\Omega(\log^{1/4} n)}$.*

Proof Sketch: We use a construction similar to the one in Theorem 4, where instead of having only one buyer of type 1, we have $m - 1$ buyers of type 1. For each type-1 buyer j and each bundle S'_i there is a special shadow-copy $S'_{i,j}$ that is only desired by this particular buyer and has value just ϵ less than the value of S'_i . So, if there is a type 1 buyer before the type 2 buyer, then the first type-1 buyer who arrives will act just like in the one Theorem 4 – he will buy an identical bundle to the proof of Theorem 4 and will not take the shadow copy, since its value is a tiny bit smaller; all later type-1 bidders will prefer their own shadow copy to any of the original sets. \square

4. UNLIMITED SUPPLY: A LOGARITHMIC APPROXIMATION

In this section we prove a logarithmic bound for the unlimited supply case for buyers with general valuations (and we should also remark that the proof of the limited supply case builds on the proof here for a single buyer). By unlimited supply we mean that the seller is able to sell any number of units of each item, and they each have zero marginal cost to the seller. For simplicity, we assume that no buyer is interested in more than a single copy of an item, and therefore the valuation is still over subsets of J .

Unlike the previous section, in this section we make no assumption about the valuation function being subadditive. We do not even assume valuations are necessarily monotone (a monotone valuation is one such that for all $S \subseteq T$, we have $v_i(S) \leq v_i(T)$, also called the *free disposal* property), so the maximum valuation for buyer i may occur at some $S \neq J$. The only assumptions we will make are that we are given the value $H = \max_{i,S} v_i(S)$ (though we will relax this later) and that the empty set has zero value to all buyers, i.e. $v_i(\emptyset) = 0$.

We prove that a random single price achieves an $O(\log m + \log n)$ approximation to the social optimum for buyers with general valuation functions. Recall, that such an approximation im-

mediately implies an approximation for the maximum revenue the seller can extract from the buyers.

Before describing the argument, let us introduce some additional useful notation. We denote by $u_{i,p}$ the utility of buyer i when the single posted price is p , and by $S_{i,p}$ the set of items that maximizes its utility, i.e., $S_{i,p} = \text{DemandPrices}(i, p, J)$.

The following lemma states that by decreasing the single posted price, the seller never sells fewer items. (For a proof see appendix A.)

LEMMA 6. *Let $p, p' \in \mathbb{R}$ such that $p > p' \geq 0$. Then, for every buyer i , $|S_{i,p'}| \geq |S_{i,p}|$.*

For clarity we start by presenting the case of a single buyer $i \in \{1, \dots, n\}$. Recall that $H = H_i = \max_S \mathbf{v}_i(S)$ is the maximum valuation of buyer i . Clearly, the profit we can extract from buyer i is at most H . Recall that $u_{i,p}$ is the maximum utility the buyer can achieve for a single posted price p , i.e., $u_{i,p} = \max_{S \subseteq J} \mathbf{u}_{i,p}(S)$.

Let $\mathbf{F}(p)$ denote the number of items our buyer purchases under the fixed single price p .⁵ We will analyze the *demand curve* which is defined as follows: the horizontal axis measures the “market price” p (the price we set on all of the items) and the vertical axis measures the number of items $\mathbf{F}(p)$ the buyer purchases at this price. Note that Lemma 6 implies that the function $\mathbf{F}(p)$ is monotonically non-increasing, as in Figure 1.

Let $p_0 = 0 < p_1 < \dots < p_L \leq H$ be such that $\mathbf{F}(p_l) = \mathbf{F}(p)$ for all $p \in [p_l, p_{l+1})$ and $\mathbf{F}(p_l) < \mathbf{F}(p_{l+1})$, for all l ; in other words, p_0, \dots, p_L are all the relevant (transition) points on the demand curve. Let us denote by $n_l = \mathbf{F}(p_l)$, for all l . (Note that since the number of items decreases with each p_l , we have that $L \leq n$.)




Figure 1: The demand curve. The horizontal axis measures the “market price” and the vertical axis measures how many items the buyer will buy at each given market price.

We will prove next a fact which is essential to our analysis, namely that the maximum valuation H of our buyer (which is also the maximum revenue our seller can extract from the buyer) is exactly the area under the \mathbf{F} –curve. Formally:

⁵In order to ensure that \mathbf{F} well defined, when there are ties we assume for simplicity that the buyer purchases the smallest bundle of maximum utility. So, $\mathbf{F}(p) = \min\{|S| : \mathbf{u}_{i,p}(S) = u_{i,p}\}$. This is the worst case for revenue to the mechanism.

LEMMA 7.

$$\max_S \mathbf{v}_i(S) = H = \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l).$$

PROOF. When the price increases from p_0 to p_1 our buyer switches from buying n_0 items to buying n_1 items, and is exactly indifferent at price p_1 . Since $u_{i,p_j} = \mathbf{v}_i(S_{i,p_j}) - n_j p_j$, this means we have $u_{i,p_1} = \mathbf{u}_{i,p_1}(S_{i,p_0}) = u_{i,p_0} - n_0 \cdot (p_1 - p_0)$. In general, for every $l > 1$, since at price p_l our buyer switches from purchasing n_{l-1} items to purchasing n_l items (and is indifferent between the two sets), we have

$$u_{i,p_l} = u_{i,p_{l-1}} - n_{l-1} \cdot (p_l - p_{l-1}).$$

So, summing all these up we obtain the desired result, $H = u_{i,p_0} -$

$$u_{i,p_L} = \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l). \quad \square$$

Let us define $q_l = \frac{H}{2^{l-1}}$, for $l \geq 1, l \in \mathbb{Z}$.⁶ We can prove (see Appendix A) that for any s , the area under \mathbf{F} is bounded above by $O(\sum_{l=1}^s q_l \cdot \mathbf{F}(q_l) + nH/2^s)$. Formally:

LEMMA 8. *For any $s \geq 1$,*

$$H = \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l) \leq 2 \cdot \sum_{l=1}^s q_l \cdot \mathbf{F}(q_l) + n \frac{H}{2^s}.$$

Let $s = \log(2n)$. Combining Lemma 7 and Lemma 8, we obtain:

$$H \leq 4 \cdot \sum_{l=1}^s q_l \cdot \mathbf{F}(q_l). \quad (1)$$

Since $q_l \cdot \mathbf{F}(q_l)$ represents the revenue obtained by the single posted price q_l , this implies that there exists a single posted price $p \in \{q_l | l \in \{1, \dots, s\}\}$ which gives an $O(\log n)$ -approximation for H .

Inequality (1) also implies that *RANDOM Single Price* is $4 \log(2n)$ -competitive with respect to the social optimum.

We can extend the analysis to multiple buyers in a direct way. The only point to notice is that we are given only $H = \max_i H_i = \max_{i,S} \mathbf{v}_i(S)$ and not the individuals $H_i = \max_S \mathbf{v}_i(S)$. Thus we need to run *RANDOM Single Price* with $s = \lceil \log(2nm) \rceil$ and guarantee an approximation of $O(\log(n) + \log(m))$. Specifically, in Appendix A we prove,

THEOREM 9. *In the case of a single buyer ($m = 1$), the RANDOM Single Price Mechanism guarantees a $4 \log(2n)$ approximation with respect to the social optimum. For any number of buyers m it guarantees $O(\log(n) + \log(m))$ -approximation with respect to the social optimum.*

Lower bound, single buyer, multiple prices: We now give a simple lower bound that holds for a general posted price algorithm that may have a different price for every item, even for the case of just one buyer. Consider a single buyer whose valuation for a set S is $\mathbf{v}(S) = \sum_{i=1}^{|S|} \frac{1}{i}$. The seller has to set a price p_j for each item j . Given his valuation functions, even though his total valuation is $\Omega(\log n)$, the buyer will never spend more than 1. In particular, given the prices, the optimal strategy of the buyer is to sort the items in increasing price order, and then buy a prefix of size k such that each item will have a cost at most $1/k$. Hence the revenue is at most 1. This establishes:

THEOREM 10. *There is a single submodular valuation function for which the revenue of any posted price mechanism has approximation ratio $\Omega(\log n)$ with respect to the social optimum.*

⁶Note that p_j increases with j while q_l decreases with l .

5. LIMITED SUPPLY: SPECIAL CASES

We now return to the case of limited supply and present two interesting classes of valuations for which a random single price does in fact achieve a logarithmic approximation.

5.1 Simple Submodular Valuations

We show here that if we have buyers with a subclass of submodular valuations that we call *simple submodular* valuations, then a random single price *does* achieve revenue within a logarithmic factor of the optimum social welfare. We say that \mathbf{v} is simple submodular if it is submodular and if for all sets S , the last item x in the greedy ordering of S satisfies $\mathbf{v}(x \mid S - \{x\}) \leq \mathbf{v}(x \mid T)$ for all $T \not\ni x$ such that $|T| \leq |S| - 1$. Here, the “greedy ordering” is the ordering in which one first chooses the item x_1 in S of highest individual valuation, then the item x_2 in S of highest marginal valuation given x_1 and so on. This class generalizes the unit-demand case (where for $|S| \geq 2$ the last item x in the greedy ordering satisfies $\mathbf{v}(x \mid S - x) = 0$) as well as the additive case and the submodular symmetric (multi-unit) setting [23].

THEOREM 11. *Consider $s = \lfloor \log_2(2n^2) \rfloor$. If all the buyers have simple submodular valuations, then the RANDOM Single Price mechanism achieves revenue within an $O(\log(n) + \log(m))$ factor of the optimum social welfare.*

Proof Sketch: Lehmann et al. [23] show that for submodular valuations, if one gives items to buyers in the order of maximum marginal valuation (i.e., at each step choose item x and buyer i to maximize $\mathbf{v}_i(x \mid S_i)$ where S_i is the set of items already given to buyer i), then this greedy procedure produces an allocation with social welfare within a constant factor of optimal. Let us call this the LLN allocation.

Note that by submodularity, these marginal values are non-increasing. Imagine that we halt this process once the marginal valuations drop below p , and let $\mathbf{F}(p)$ denote the number of items allocated by that point. As in the proof of Theorem 9, the expected value of $p \cdot \mathbf{F}(p)$ is within a logarithmic factor of the social welfare of the overall LLN allocation. Thus, it suffices to show that allowing buyers to enter in an arbitrary order and purchase at price p achieves revenue within a constant factor of $p \cdot \mathbf{F}(p)$.

In particular, suppose that buyer j is assigned k_j items of marginal valuation at least p in the LLN allocation. When buyer j arrives, the claim is that he will either purchase all remaining items in the *second half* of his assigned bundle or else purchase at least $k_j/2$ items total. In particular, suppose not and let x be an unpurchased item in the second half, let T be items purchased, and let S'_j be the items in his assigned bundle up through item x . By definition of the LLN procedure we have $\mathbf{v}_j(x \mid S'_j - x) > p$ and by definition of simple submodular valuations we have $\mathbf{v}_j(x \mid T) \geq \mathbf{v}_j(x \mid S'_j - x)$ since $|T| < |S'_j|$. This in turn then implies by a standard matching argument that at least $\mathbf{F}(p)/4$ items are sold in total, as desired. \square

5.2 The Multi-Unit Case

We consider here the multi-unit setting [11], and we show how the ideas from the unlimited supply setting (Section 4) can be applied to get a logarithmic approximation in this case as well. The formal setting here is the following. Assume that we have only one item for sale, but n copies of it. We also have m customers, and each customer i has a valuation function $v_i : \{1, \dots, n\} \rightarrow \mathbb{R}$, where $v_i(q)$ encodes his value for obtaining q items, and these valuation functions can be arbitrarily complicated. The setting is essentially the symmetric valuations case in combinatorial auctions as described in Lehmann et al. [23] or in Vickery [27]; note however that we are interested in general symmetric valuations, not only the submodular case.

We first point out that if we make *no* assumptions, then no single price can guarantee an $o(n)$ approximation to the social welfare. In particular, suppose that the first customer has value 1 on any set of size 1, and the second customer has value n on the entire set of n items but value 0 on any other set. In this case, the maximum possible revenue is 1. However, note that this case is a bit peculiar because one buyer gets all the items in the optimal allocation. We show here that this is the only barrier to good revenue. In particular, we show that under the assumption that the optimal allocation gives at most a $(1 - \epsilon)$ fraction of the items to any one buyer, our single pricing scheme can achieve a logarithmic approximation in this setting as well. Specifically:

THEOREM 12. *For the case of n identical items, under the assumption that the optimal (or a near-optimal) allocation gives at most a $(1 - \epsilon)$ fraction of the items to any one buyer, the RANDOM Single Price Mechanism guarantees an $O(\frac{1}{\epsilon} \log(n) + \log(m))$ -approximation with respect to the social optimum.*

Proof Sketch: Let n_i be the number of items given to buyer i in the optimal allocation (the one maximizing social welfare), and let v_i be the value of this set to buyer i (so $v_i = v_i(n_i)$). By assumption, all $n_i \leq (1 - \epsilon)n$. Now, for each buyer i , imagine drawing the monotone demand curve $\mathbf{F}_i(p)$ of number of items desired versus price, but where we cap the number of items at n_i . That is, $\mathbf{F}_i(p)$ is the number of items buyer i would purchase in a store with n_i items all at price p . So, $\mathbf{F}_i(0) \leq n_i$ and it eventually drops to 0. The integral of this function is exactly v_i . If we sum up these curves, we get a global curve $\mathbf{F} = \mathbf{F}_1 + \mathbf{F}_2 + \dots + \mathbf{F}_m$ whose integral is the optimum social welfare.

Now, let p be a price whose rectangle has large area $p \cdot \mathbf{F}(p)$. The claim is that if bidders arrive in an arbitrary order, the seller makes at least an ϵ fraction of this amount $\epsilon p \cdot \mathbf{F}(p)$ by pricing items at p . The reason is that by definition of \mathbf{F} , when a buyer i arrives, he will purchase at least $\mathbf{F}_i(p)$ items if that many are available. Since the only way it is possible for this number of items to not be available is for at least ϵn to have been already sold, this means the total number of items purchased is at least $\min(\mathbf{F}(p), \epsilon n) \geq \epsilon \mathbf{F}(p)$. This then implies the desired result. \square

We can also extend the above argument to the case that the seller has n copies of k distinct items, so long as the optimal allocation gives no buyer more than $(1 - \epsilon)n$ items total.

THEOREM 13. *For the case of n copies of k distinct items, under the assumption that the optimal allocation gives at most $(1 - \epsilon)n$ items to any one buyer, the RANDOM Single Price Mechanism guarantees an $O(\frac{k}{\epsilon}(\log(n) + \log(m)))$ -approximation with respect to the social optimum.*

Proof Sketch: The same argument as above applies, where n_i is the *total* number of items given to buyer i in the optimal allocation. The difference is that $\mathbf{F}(p)$ could now be as large as nk , so $\min(\mathbf{F}(p), \epsilon n) \geq \epsilon \mathbf{F}(p)/k$. \square

6. CONCLUSIONS AND OPEN PROBLEMS

We show that single posted price mechanisms are surprisingly powerful, achieving revenue within a logarithmic factor of the total social welfare for unlimited supply settings for buyers with general valuation functions (not just single-minded or unit-demand) and achieving a $2^{O(\sqrt{\log n \log \log n})}$ approximation for the limited supply case with subadditive buyers. These are the best revenue guarantees known for *any* item-pricing scheme, and in the unlimited supply setting match the best possible guarantee by any item pricing scheme.

We also provide a $2^{\Omega(\log^{1/4} n)}$ lower bound on the revenue of any single posted price mechanism for subadditive buyers (even XOS buyers) in the limited supply setting, showing that even for buyers with known valuation functions, the gap still exists. Since social welfare *does* have a logarithmic approximation using a single price [12, 10], this demonstrates a clear distinction between revenue maximization and social welfare maximization in the limited supply setting.

Note that our lower bound for limited supply does not apply if one allows the seller to use different prices on different items. An interesting open question is whether an improved upper bound is possible using multiple prices, or on the other hand whether an alternative lower bound can be given for that case. In particular, it is an open question if the lower bound can be extended even to the case where the seller is allowed to use just *two* prices. A second open question is whether our lower bound (which uses XOS buyers) can be extended to the more restricted class of submodular buyers, or whether alternatively a polylog(n) upper bound can be obtained in that case.

Acknowledgments: We thank Shahar Dobzinski and Jason D. Hartline for a number of useful discussions.

7. REFERENCES

[1] G. Aggarwal, T. Feder, R. Motwani, and A. Zhu. Algorithms for multi-product pricing. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*, pages 72–83, 2004.

[2] M.-F. Balcan and A. Blum. Approximation Algorithms and Online Mechanisms for Item Pricing. In *ACM Conference on Electronic Commerce*, 2006.

[3] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 605 – 614, 2005.

[4] M.-F. Balcan, A. Blum, and Y. Mansour. Single price mechanisms for revenue maximization in unlimited supply combinatorial auctions. Technical Report CMU-CS-07-111, Carnegie Mellon University, February 2007.

[5] L. Blumrosen and N. Nisan. Combinatorial auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[6] P. Briest, M. Hoefer, and P. Krysta. Stackelberg network pricing games. In *Proceedings of the 25th International Symposium on Theoretical Aspects of Computer Science (STACS)*, 2008.

[7] P. Briest and P. Krysta. Single-Minded Unlimited Supply Pricing on Sparse Instances. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[8] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. Springer-Verlag, 2005.

[9] E. D. Demaine, U. Feige, M. Hajiaghayi, and M. R. Salavatipour. Combination Can Be Hard: Approximability of the Unique Coverage Problem. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 162 – 171, 2006.

[10] S. Dobzinski. Two Randomized Mechanisms for Combinatorial Auctions. In *APPROX*, 2007.

[11] S. Dobzinski and N. Nisan. Mechanisms for Multi-Unit Auctions. In *EC*, 2007.

[12] S. Dobzinski, N. Nisan, and M. Schapira. Truthful Randomized Mechanisms for Combinatorial Auctions. In

Proceedings of the thirty-eighth annual ACM symposium on Theory of computing, pages 644–652, 2006.

[13] K. Elbassioni, R. Sitters, and Y. Zhang. A Quasi-PTAS for Profit-Maximizing Pricing on Line Graphs. In *ESA*, 2007.

[14] U. Feige. On maximizing Welfare when Utility Functions are Subadditive. In *STOC*, 2006.

[15] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proceedings of the 34th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2002.

[16] A. Goldberg, J. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 735–744, 2001.

[17] A. Grigoriev, J. van Loon, R. Sitters, and M. Uetz. How to Sell a Graph: Guidelines for Graph Retailers. Meteor Research Memorandum RM/06/001, Maastricht University, 2005.

[18] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On Profit-Maximizing Envy-Free Pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1164 – 1173, 2005.

[19] J. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[20] J. Hartline and V. Koltun. Near-Optimal Pricing in Near-Linear Time. In *9th Workshop on Algorithms and Data Structures*, pages 422–431, 2005.

[21] R. Lavi. Combinatorial auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[22] R. Lavi and C. Swamy. Truthful and Near-optimal Mechanism Design via Linear Programming. In *46th Annual IEEE Symposium on Foundations of Computer Science*, 2005.

[23] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 2006.

[24] D. Lehmann, L. I. Ocallaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)*, 49:577 – 602, 2002.

[25] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.

[26] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[27] W. Vickery. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, 1961.

APPENDIX

A. PROOFS

Lemma 1 Assume that \mathbf{v}_i are all subadditive. Let T_1, \dots, T_m be an arbitrary feasible allocation and let $\alpha = \frac{1}{4 \log(2n^2)}$. There exists p and subsets $L_i \subseteq T_i$ such that L_1, \dots, L_m is an allocation supported at price p and furthermore

$$\sum_{i=1}^m p|L_i| \geq \alpha \sum_{i=1}^m \mathbf{v}_i(T_i).$$

PROOF. Consider L_i to be the subset of T_i that buyer i would purchase if he were in a store where everything is priced at p and he is only allowed to see the items in T_i . The fact that $\sum_{i=1}^m p|L_i| \geq \alpha \sum_{i=1}^m \mathbf{v}_i(T_i)$ then follows from Theorem 9. The fact L_1, \dots, L_m is an allocation supported at price p follows from properties of subadditive valuation functions. Specifically, consider an arbitrary subset W_i of L_i . Since L_i is the set that buyer i would purchase with items in T_i priced at p and items in $J \setminus T_i$ removed, we know that

$$\mathbf{v}_i(L_i) - \mathbf{v}_i(L_i \setminus W_i) \geq p \cdot |W_i|.$$

(If the above were not the case, the buyer would have purchased the set $L_i \setminus W_i$ instead of L_i .) In addition, since \mathbf{v}_i is subadditive we also have:

$$\mathbf{v}_i(L_i) \leq \mathbf{v}_i(W_i) + \mathbf{v}_i(L_i \setminus W_i).$$

Combining these together we get that $\mathbf{v}_i(W_i) \geq p \cdot |W_i|$ as desired. \square

Lemma 2 Let L_1, \dots, L_m be an allocation supported at price p . Let S_1, \dots, S_m be the allocation produced by Generate Allocation when run with the price parameter $p/2$. Then:

$$\sum_{i=1}^m \mathbf{v}_i(S_i) \geq \sum_{i=1}^m (p/2)|L_i|.$$

PROOF. Define $\text{ALG}(i, S_1, \dots, S_{i-1})$ as the total valuation that our algorithm acquires from the set of bidders i, \dots, m on the set of items $J \setminus (S_1 \cup S_2 \dots \cup S_{i-1})$. Let $\text{OPT}_{rev}(i, L_1, \dots, L_m, S_1, \dots, S_{i-1})$ be the total revenue obtained at price $p/2$ from the set of bidders i, \dots, m if buyer $j \geq i$ would be allocated set $L_j \setminus (S_1 \cup S_2 \dots \cup S_{i-1})$. We clearly have

$$\text{ALG}(i, S_1, \dots, S_{i-1}) = \text{ALG}(i+1, S_1, \dots, S_i) + \mathbf{v}_i(S_i)$$

and also

$$\begin{aligned} \text{OPT}_{rev}(i, L_1, \dots, L_m, S_1, \dots, S_{i-1}) \\ \leq \text{OPT}_{rev}(i+1, L_1, \dots, L_m, S_1, \dots, S_i) + \\ (p/2)|L_i \setminus (S_1 \cup S_2 \dots \cup S_{i-1})| + (p/2)|S_i|. \end{aligned}$$

The second inequality follows from the fact that we lose at most $(p/2)|L_i \setminus (S_1 \cup S_2 \dots \cup S_{i-1})|$ from using up buyer i , and at most $(p/2)|S_i|$ from using up S_i since S_i might contain items that are in $L_j \setminus (S_1 \cup S_2 \dots \cup S_{i-1})$ for $j > i$. But

$$\begin{aligned} & |L_i \setminus (S_1 \cup S_2 \dots \cup S_{i-1})|(p/2) \\ & \leq \mathbf{v}_i(L_i \setminus (S_1 \cup S_2 \dots \cup S_{i-1})) \\ & \quad - |L_i \setminus (S_1 \cup S_2 \dots \cup S_{i-1})|(p/2) \\ & \leq \mathbf{v}_i(S_i) - (p/2)|S_i|. \end{aligned}$$

The first inequality follows from the fact that the allocation L_1, \dots, L_m is supported at price p and the second one from the fact that buyer i prefers S_i to the set $L_i \setminus (S_1 \cup S_2 \dots \cup S_{i-1})$. So

$$\mathbf{v}_i(S_i) \geq (p/2)|S_i| + |L_i \setminus (S_1 \cup S_2 \dots \cup S_{i-1})| \cdot (p/2),$$

and therefore

$$\begin{aligned} \text{ALG}(i, T_1, \dots, T_m, S_1, \dots, S_{i-1}) \\ = \text{ALG}(i+1, T_1, \dots, T_m, S_1, \dots, S_i) + \mathbf{v}_i(S_i). \end{aligned}$$

and

$$\begin{aligned} \text{OPT}_{rev}(i, T_1, \dots, T_m, S_1, \dots, S_{i-1}) \\ \leq \text{OPT}_{rev}(i+1, T_1, \dots, T_m, S_1, \dots, S_{i-1}) + \mathbf{v}_i(S_i). \end{aligned}$$

We also have $\text{OPT}_{rev}(1, L_1, \dots, L_m) = \sum_{i=1}^m (p/2)|L_i|$, and all these imply:

$$\sum_{i=1}^m \mathbf{v}_i(S_i) \geq \sum_{i=1}^m (p/2)|L_i|,$$

as desired. \square

Lemma 6 Let $p, p' \in \mathbb{R}$ such that $p > p'$. Then, for every buyer i , $|S_{i,p'}| \geq |S_{i,p}|$.

PROOF. Let us fix a buyer i . Assume that $|S_{i,p}| = k$. By definition, since $S_{i,p}$ is a set of items that maximizes the buyer's utility under the pricing vector $\vec{p} = p\mathbf{1}_n$, for all subsets $T \subseteq J$ we have:

$$\mathbf{u}_{i,p}(S_{i,p}) = \mathbf{v}_i(S_{i,p}) - p \cdot k \geq \mathbf{u}_{i,p}(T) = \mathbf{v}_i(T) - p \cdot |T|.$$

Assume now that $p = p' + \epsilon$, $\epsilon > 0$, and let T be an arbitrary subset of J with $|T| = k'$, $k' < k$. Then we clearly have:

$$\begin{aligned} \mathbf{v}_i(S_{i,p}) - p' \cdot k &= \mathbf{v}_i(S_{i,p}) - (p - \epsilon) \cdot k \\ &= \mathbf{v}_i(S_{i,p}) - p \cdot k + k' \cdot \epsilon \\ &> \mathbf{v}_i(T) - p \cdot k' + k' \cdot \epsilon \\ &\geq \mathbf{v}_i(T) - (p - \epsilon) \cdot k' \\ &= \mathbf{v}_i(T) - p' \cdot |T|. \end{aligned}$$

Therefore, for all subsets $T \subseteq J$ with $|T| = k'$, $k' < k$, we have:

$$\mathbf{u}_{i,p'}(S_{i,p}) = \mathbf{v}_i(S_{i,p}) - p' \cdot k > \mathbf{v}_i(T) - p' \cdot |T| = \mathbf{u}_{i,p'}(T).$$

This then implies that any set of items $S_{i,p'}$ that maximizes buyer's i utility under the pricing vector \vec{p}' satisfies $|S_{i,p'}| \geq |S_{i,p}|$, as desired. \square

Lemma 8 For any $s \geq 1$,

$$H = \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l) \leq 2 \cdot \sum_{l=1}^s q_l \cdot \mathbf{F}(q_l) + n \frac{H}{2^s}.$$

PROOF. By Lemma 7 we have $H = \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l)$. Now we can bound the sum as follows,

$$\begin{aligned} & \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l) \\ &= \sum_{l=1}^L (n_{l-1} - n_l) \cdot p_l \\ &= \sum_{l:p_l \geq \frac{H}{2^s}} (n_{l-1} - n_l) \cdot p_l + \sum_{l:p_l < \frac{H}{2^s}} (n_{l-1} - n_l) \cdot p_l \\ &\leq \sum_{l:p_l \geq \frac{H}{2^s}} (n_{l-1} - n_l) \cdot p_l + n \frac{H}{2^s} \end{aligned}$$

Consider the prices that fall in the range $[q_l, q_{l-1})$, and assume they are $p_j \leq \dots \leq p_{j+k}$. Clearly we have that each price in the range is at most q_{l-1} . Since $\sum_{b=j}^{j+k} (n_{b-1} - n_b) = n_{j-1} - n_{j+k} \leq \mathbf{F}(q_l)$, we have,

$$H \leq \sum_{l=1}^s q_{l-1} \cdot \mathbf{F}(q_l) + n \frac{H}{2^s} = 2 \sum_{l=1}^s q_l \cdot \mathbf{F}(q_l) + n \frac{H}{2^s},$$

as desired. \square

Theorem 9 In the case of a single buyer ($m = 1$), the *RANDOM Single Price* Mechanism guarantees a $4 \log(2n)$ approximation with respect to the social optimum. For any number of buyers m it guarantees $O(\log(n) + \log(m))$ -approximation with respect to the social optimum.

PROOF. For $m = 1$, the desired competitive ratio follows from (1) and from the fact that the expected profit of our mechanism is $\frac{1}{s} \sum_{l=1}^s q_l \cdot \mathbf{F}(q_l)$.

Assume $m \geq 1$. Let $H_i = \max_S(\mathbf{v}_i(S))$, and let \mathbf{F}_i be the curve corresponding to buyer i ; so $H = \max_i H_i$. Let $s_i = s - \log(H/H_i)$, which is the effective index for the i -th buyer.⁷ By Lemma 8, applied to buyer i , we have,

$$H_i \leq 2 \cdot \sum_{l=1}^s q_l \cdot \mathbf{F}_i(q_l) + n \frac{H_i}{2^{s_i}} = 2 \cdot \sum_{l=1}^s q_l \cdot \mathbf{F}_i(q_l) + n \frac{H}{2^s}. \quad (2)$$

Summing over all the buyers we have:

$$\sum_{i=1}^m H_i \leq 2 \cdot \sum_{i=1}^m \sum_{l=1}^s q_l \cdot \mathbf{F}_i(q_l) + nm \frac{H}{2^s}. \quad (3)$$

Since $\sum_{i=1}^m H_i \geq H$, $s = \log(2nm)$, combining (3) together with the fact that the expected profit of our mechanism is

$$\sum_{i=1}^m \left(\frac{1}{s} \sum_{l=1}^s q_l \cdot \mathbf{F}_i(q_l) \right),$$

we get an approximation ratio of $4 \log(2nm) = O(\log n + \log m)$. \square

Note: Note that the $O(\log(m))$ factor is attributed directly to the variation in H_i . Assume that $H_i = H$ for every buyer i . Then $\sum_{i=1}^m H_i = mH$ and it is sufficient to set $s = \log(2n)$.

Removing the assumption of known H : We have assumed so far that the maximum valuation H (over all buyers and all bundles) is known to the mechanism. We can remove this assumptions using rather “standard tricks”. One immediate generalization is that if instead we are just given an upper-bound H' on H , with the guarantee that $H' \leq \alpha H$ for some given value α , then the mechanism *RANDOM Single Price* is an $O(\log n + \log m + \log \alpha)$ approximation. In particular, this implies that if we are simply given a polynomial upper bound H' on H , i.e., $H' \leq \text{poly}(m, n) \times H$, then we still get an $O(\log n + \log m)$ bound.

Alternatively, if we have no upper bound on H at all, but we assume at least that $H \geq 1$, then select H' at random from the probability distribution where $\Pr[H' = 2^i] = \frac{c}{i \log^2 i}$, for some constant $c > 0$. Now we can run mechanism *RANDOM Single Price* with the selected H' and the parameter s . The probability that *RANDOM Single Price* selects a given price $p = 2^k$ is $\frac{1}{s} \sum_{i=k}^{s+k} \Pr[H' = 2^i] \geq \frac{c}{(s+k) \log^2(s+k)}$. This implies that the approximation ratio, for $s = \log(2mn)$ and $p \leq H$ is $O(\log(nmH) \log^2(\log(mnH)))$.

B. ADDITIONAL RELATED WORK

As noted in the Introduction, the setting we analyze is related to the Combinatorial Auctions setting. In the Computer Science community, there have recently been two main threads of work in the

⁷For simplicity we assume that $\log(H/H_i)$ is an integer.

context of Combinatorial Auctions: designing good algorithms [18, 1, 7, 2, 9, 20] and mechanisms [15, 16, 3] for *revenue maximization* in the unlimited supply setting [19], and designing algorithms [14] and computationally efficient mechanisms [24, 22, 12, 26, 10] to optimize *social welfare* in the limited supply setting.

Substantial effort has been devoted to find a computationally efficient combinatorial auction which approximates the social welfare well [24, 22, 12, 26, 10]. On the other hand, as pointed out in [26], much less is known about designing combinatorial auctions that maximize the auctioneer’s *revenue*. In particular, as opposed to the social welfare goal, where obtaining a truthful mechanism is easy ignoring the computational constraints (due to the celebrated VCG mechanism), for the revenue maximization goal no such mechanism is known. The most notable positive result so far for revenue maximization in the limited supply setting is due to Dobzinski, Nisan and Schapira [12], who present a simple random-sampling based truthful mechanism, that provides an $O(\sqrt{n})$ -approximation for bidders with general valuation functions, *both* for the social welfare and revenue maximization objectives. They additionally show an $O(\log(n))$ -approximation to *social welfare* for the special case of XOS bidders, recently generalized to the somewhat larger class of subadditive bidders by [10]. At the heart of this mechanism is an item-pricing in which all items get the same random price and then bidders enter one at a time and purchase what they want most at that price, precisely the mechanism we have analyzed throughout this paper. Dobzinski (personal communication) points out that if one allows more general bundle-pricing mechanisms, then the results of [10] can be adapted to provide polylogarithmic revenue guarantees for subadditive buyers: the mechanism essentially avoids the lower bound of Theorem 4 by refusing to sell small bundles to any buyer. Hartline (personal communication) points out that if furthermore computational efficiency is not an issue, then VCG with a random offset will achieve logarithmic revenue guarantees for general buyers in the limited supply case. These mechanisms, however, require pricing *bundles*, and so do not apply to the setting in which objects must be priced per item, which includes the vast majority of sales in the world today.

Revenue maximization in the *unlimited supply* setting has also become increasingly popular in the past few years; for a recent survey see [19]. Note that much of the work on revenue maximization in combinatorial auctions has focused on *item pricing*, in part because of its wide applicability. Some of these results are truthful mechanisms and some are not [1, 2, 7, 17, 18, 13, 20], though Balcan et al. [3] give a generic reduction to convert any (non-truthful) item-pricing to a truthful mechanism when the number of bidders is sufficiently large as a function of various measure of complexity of the class of item pricings. Guruswami et al. [18] give an $O(\log m + \log n)$ -approximation both for the case of single minded and unit-demand bidders; furthermore, Demaine, Feige, Hajiaghayi, and Salavatipour [9] show that it is hard to approximate the maximum revenue within a factor of $\log^\delta n$, for some $\delta > 0$, assuming that $\text{NP} \not\subseteq \text{BPTIME}(2^{n^\epsilon})$ for some $\epsilon > 0$, even for the case of single-minded bidders.

In the context of Bayesian mechanism design, Myerson derives in a seminal paper [25] the optimal auction for selling a single item given that the bidders’ true valuations for the item come from some known *prior distribution*. His mechanism generalizes trivially to any single-parameter agent setting with arbitrary supply constraints or costs to the auctioneer for the outcome produced. However, no such characterization is known for the more general case of multi-parameter settings.