Single Price Mechanisms for Revenue Maximization in Unlimited Supply Combinatorial Auctions

Maria-Florina Balcan* **Avrim Blum*** **Yishay Mansour†**

February 2007
CMU-CS-07-111

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

* School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. {ninamf,avrim}@cs.cmu.edu
† School of computer Science, Tel-Aviv University. mansour@cs.tau.ac.il.

Abstract

In this paper we generalize a result of Guruswami et. al [7], showing that in unlimited-supply combinatorial auctions, a surprisingly simple mechanism that offers the same price for each item achieves revenue within a logarithmic factor of the total social welfare for bidders with general valuation functions (not just single-minded or unit-demand bidders as in [7]). We also extend this to the limited-supply setting for the special case of bidders with additive valuations. These are both settings for which Likhodedov and Sandholm [9] provide logarithmic approximations but via much more complex bundle-pricing mechanisms.

Research supported in part by NSF grants CCR-0105488, CCR-0122581, and IIS-0121678.

Keywords: Mechanism Design, Profit Maximization, Unlimited Supply, Combinatorial Auctions

1 Introduction

Consider a retailer with n goods for sale, each in unlimited supply, and m customers with unknown and possibly highly complex valuation functions over these goods. The retailer would like to assign prices to the goods to achieve revenue close to the total social welfare (the maximum revenue achievable if it were possible to sell to each customer the bundle she values the most at exactly her valuation for it). Guruswami et. al [7] show that in such a setting, namely an unlimited supply combinatorial auction, if customers are unit-demand or single-minded then 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). In this paper we generalize the result of [7] by showing that the restriction to single-minded or unit-demand valuations is not required, and that the result in fact holds for general (not even necessarily monotone) valuation functions. Note that Likhodedov and Sandholm [9] give a more complex mechanism, based on virtual valuations, that achieves a logarithmic approximation for bidders with general monotone valuation functions; however, the resulting mechanism requires pricing bundles and is not necessarily a pricing of the individual items.

Related work: Revenue maximization auctions and algorithmic pricing problems have generated a great deal of interest recently; for a comprehensive survey see the chapter "Profit Maximization in Mechanism Design" of Hartline and Karlin in [10]. Note that much of the previous algorithmic and mechanism design work concerning combinatorial auctions in the context of revenue maximization has focused on item pricing [1, 4, 3, 5, 7, 8].

In a *limited supply* setting (one copy of each item), Dobzinski, Nisan and Schapira [6] present a $O(\sqrt{n})$-approximation for the general case, and a $O(\log^3(n))$-approximation for the special case of submodular bidders. We also note that Awerbuch, Azar and Meyerson [2] present a logarithmic competitive mechanism for the case when the supply is limited, but large.

2 Notation and Definitions

We consider the following setting. We have a set J of n items and a set B of m bidders or customers. We are in an *unlimited supply* setting, which means that the seller is able to sell any number of units of each item, and they each have zero marginal cost to the seller.[1]

Each customer i 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 that customer. This function could be quite complex (since there are 2^n possible bundles). We will not assume valuations are necessarily monotone (a monotone valuation is one such that for all $S \subseteq T$, we have $\mathbf{v}_i(S) \leq \mathbf{v}_i(T)$, also called the *free disposal* property), so the maximum valuation for bidder i may occur at some $S \neq J$. The only assumptions we will make are that we are given the value $H = \max_{i,S} \mathbf{v}_i(S)$, though we will relax this later, and that the empty set has zero value to all bidders, i.e. $\mathbf{v}_i(\emptyset) = 0$. A pricing scheme can be described by a vector $\mathbf{p}$ of prices on all possible subsets of J. Given a price vector $\mathbf{p}$, the utility that bidder i derives from bundle S is $\mathbf{u}_{i,\mathbf{p}}(S) = \mathbf{v}_i(S) - \mathbf{p}(S)$, and this measures the bidder's satisfaction of having bought the bundle S at the price $\mathbf{p}(S)$. We define the revenue achieved by a price vector $\mathbf{p}$ to be the revenue obtained if each bidder buys the bundle (which could be empty) that maximizes her/his utility, i.e. the one with largest positive gap between its value to the bidder and its total cost. So, any mechanism that chooses prices in a manner that is independent of the valuations

[1] Or if the items have some fixed marginal cost, we have subtracted that from all valuations and the seller may not sell any item below cost.

will by definition be incentive-compatible. We denote by $S_{i,\mathbf{p}}$ the set of items that maximizes $\mathbf{u}_{i,\mathbf{p}}(S)$, i.e., $S_{i,\mathbf{p}} = \text{argmax}_{S \subseteq J}\, \mathbf{u}_{i,\mathbf{p}}(S)$.

We will consider pricing schemes of special types called *item pricings* and *single fixed pricings*. An item pricing is a pricing scheme that assigns a separate price to each item, such that the price of a bundle is just the sum of the prices of the items in it. So, an item pricing $\vec{p}$ is a price vector $\vec{p} = (p_1, \ldots, p_n) \in \mathbb{R}^+$ and the prices it induces on each bundle $S \subseteq J$ are $\vec{p}(S) = \sum_{j \in S} p_j$. A *single fixed* pricing scheme is an item pricing that assigns *the same price* to each item. So, a single fixed pricing is an item pricing of the type $\vec{p} = p\mathbf{1}_n$, where $p \in \mathbb{R}^+$ and $\mathbf{1}_n = (1, ..., 1)$ is the n-dimensional vector with all 1-entries. For simplicity, we will denote by $u_{i,p}$ the utility function of bidder i induced by the pricing vector $\vec{p} = p\mathbf{1}_n$ and by $S_{i,p}$ the set of items that maximizes $\mathbf{u}_{i,\vec{p}}(S)$ (i.e. $u_{i,p} := \mathbf{u}_{i,\vec{p}}(S_{i,p})$ and $S_{i,p} = \arg\max_S \mathbf{u}_{i,\vec{p}}(S)$ where $\vec{p} = p\mathbf{1}_n$).

The goal of the seller is to set prices on items so as to maximize his revenue, i.e. the sum of the prices of all the bundles sold, under the assumption that each bidder purchases the bundle maximizing her utility. Our mechanisms will choose their price p independently from the valuations $\mathbf{v}_i$ and so will be incentive-compatible. In Section 4 we consider a limited supply setting, and so must also specify an allocation rule; this will be incentive-compatible as well. We emphasize that our goal is to maximize revenue with an *item pricing*. If bundle-pricing is allowed, then it is not difficult to see that for the case of unlimited supply one can in fact achieve a logarithmic approximation by choosing a random price p and then pricing all *bundles* at that price (i.e., charging an admissions fee after which all items are free, as is done with rides in an amusement park).

3 The main result: A Logarithmic Approximation

We prove in this section that a single fixed pricing scheme achieves an $O(\log m + \log n)$ approximation to the total social welfare for bidders with general valuation functions. Such an approximation immediately implies an approximation for the maximum revenue the seller can extract from the bidders.

3.1 Monotonicity of bundle sizes

We start with a simple lemma that is needed for our analysis, which states that by decreasing the price the seller never sells fewer items. Specifically:

Lemma 1 *Let $p, p' \in \mathbb{R}$ such that $p > p'$. Then, for every bidder i, $|S_{i,p'}| \geq |S_{i,p}|$.*

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

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

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:

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

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

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

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

3.2 Single Bidder Case

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

We will analyze the *demand curve* which is defined as follows: the horizontal axis measures the "market price" (the price we set on all of the items) and the vertical axis measures how many items the bidder will buy at each given market price. Note that Lemma 1 implies that this curve slopes down, as in Figure 1; in other words the function $\mathbf{F}$ that specifies how many items $\mathbf{F}(p)$ our bidder buys under a fixed single price $\vec{p} = p\mathbf{1}_n$ is monotonically non-increasing.[2]

Let $p_0 = 0$, p_1, ..., $p_L \leq H$ 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, ..., 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 $L \leq n$.)

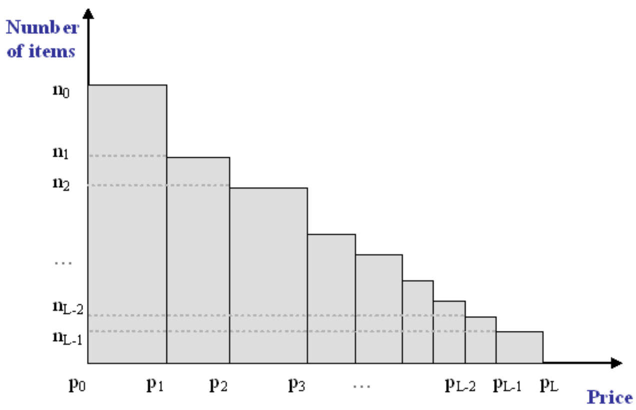

Figure 1: The demand curve. The horizontal axis measures the "market price" and the vertical axis measures how many items the bidder will buy at each given market price.

We will prove next a fact which is essential to our analysis, namely that the maximal valuation H of our bidder (which is also the maximum profit we can extract from the bidder) is at most the area under the $\mathbf{F}-$curve. Formally:

Lemma 2

$$H \leq \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l).$$

[2]In order to ensure that $\mathbf{F}$ well defined, when there are ties we assume that the bidder purchases the smallest bundle of maximum utility. So, $\mathbf{F}(p) = \min\{|S| : \mathbf{u}_{i,p}(S) = u_{i,p}\}$.

Proof: When the prices changes from $p_0\mathbf{1}_n$ to $p_1\mathbf{1}_n$ our bidder switches from buying n_0 items to buying n_1 items. Since $u_{i,p_j} = \mathbf{v}_i(S_{i,p_j}) - n_j p_j$, we have $u_{i,p_1} \geq \mathbf{u}_{i,p_1}(S_{i,p_0}) = u_{i,p_0} - n_0 \cdot (p_1 - p_0)$. In general, for every $l > 1$, since under the pricing $p_l\mathbf{1}_n$ our bidder switches from buying n_{l-1} items to buying n_l items, we have

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

So, summing all these up we obtain the desired result, namely that:

$$u_{i,p_0} - u_{i,p_L} \leq H \leq \sum_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l).$$

■

Let us denote by $\tilde{q}_l = \frac{H}{2^{l-1}}$, for $l \geq 1$, $l \in \mathbb{Z}$.[3] We will prove in the following that for any s, the area under $\mathbf{F}$ is bounded above by $O(\sum\limits_{l=1}^{s} \tilde{q}_l \cdot \mathbf{F}(\tilde{q}_l) + nH/2^s)$. Formally:

Lemma 3 *For any $s \geq 1$,*

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

Proof: By Lemma 2 we have $H \leq \sum\limits_{l=0}^{L-1} n_l \cdot (p_{l+1} - p_l)$. 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 $[\tilde{q}_l, \tilde{q}_{l-1})$, and assume they are $p_j \leq \cdots \leq p_{j+k}$. Clearly we have that each price in the range is at most $\tilde{q}_{l-1}$. Since $\sum_{b=j}^{j+k}(n_{b-1} - n_b) = n_{j-1} - n_{j+k} \leq \mathbf{F}(\tilde{q}_l)$, we have,

$$H \;\; \leq \;\; \sum_{l=1}^{s} \tilde{q}_{l-1} \cdot \mathbf{F}(\tilde{q}_l) + n\frac{H}{2^s} = 2\sum_{l=1}^{s} \tilde{q}_l \cdot \mathbf{F}(\tilde{q}_l) + n\frac{H}{2^s},$$

as desired. ■

Let $s = \log(2n)$. Combining Lemma 2 and Lemma 3, we obtain:

$$H \leq 4 \cdot \sum_{l=1}^{s} \tilde{q}_l \cdot \mathbf{F}(\tilde{q}_l). \tag{1}$$

This implies that there exists a fixed price $p \in \{\tilde{q}_l | l \in \{1, ..., s\}\}$ which gives an $O(\log n)$-approximation for H. It also implies that:

[3] Note that p_j increases with j while $\tilde{q}_l$ decreases with l.

Theorem 1 *The $RANDOM$ Mechanism (given below) with $m = 1$ bidders is truthful and $4\log(2n)$-competitive with respect to the total social welfare H.*

Proof: The truthfulness of $RANDOM$ is immediate. The desired competitive ratio follows from (1) and from the fact that the expected profit of our mechanism is $\frac{1}{s}\sum_{l=1}^{s}\tilde{q}_l \cdot \mathbf{F}(\tilde{q}_l)$. ■

Mechanism 1 $RANDOM$

Input: $m =$ number of bidders, $n =$ number of items, $H = \max_{i,S}(\mathbf{v}_i(S))$.

Step 1 Let $\tilde{q}_l = \frac{H}{2^{l-1}}$, for $l \geq 1$, $l \in \mathbb{Z}$, and let $s = \log(2nm)$

Step 2 Pick a price $\tilde{q}$ uniformly at random in $\{\tilde{q}_1, ..., \tilde{q}_s\}$ and price all items at $\tilde{q}$. I.e., output the allocation and the payment determined by $\tilde{q}\mathbf{1}_n$.

3.3 The Multiple Bidders Case

We can now extend the results in the previous section to the multiple bidders case as follows:

Theorem 2 *The $RANDOM$ Mechanism is truthful and $O(\log(n) + \log(m))$-competitive with respect to the total social welfare of the bidders.*

Proof: The truthfulness of $RANDOM$ is immediate. Let $H_i = \max_S(\mathbf{v}_i(S))$, and let $\mathbf{F}_i$ be the curve corresponding to bidder i; so $H = \max_i H_i$. Let $s_i = s - \log(H/H_i)$, which is the effective index for the i-th bidder.[4] By Lemma 3, applied to bidder i, we have,

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

Summing over all the bidders we have:

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

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 $\sum_{i=1}^{m}\left(\frac{1}{s}\sum_{l=1}^{s}\tilde{q}_l \cdot \mathbf{F}_i(\tilde{q}_l)\right)$, we get an approximation ratio of $4\log(2nm) = O(\log n + \log m)$. ■

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

[4]For simplicity we assume that $\log(H/H_i)$ is integer.

3.4 Removing the assumption of known H

We have assumed so far that the maximum valuation H (over all bidders and all bundles) is known to the mechanism. 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 we can use $s = \log(2\alpha mn)$ in mechanism $RANDOM$ and the result 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 $\mathbf{Pr}[H' = 2^i] = \frac{c}{i \log^2 i}$, for some constant $c > 0$. Now we can run mechanism $RANDOM$ with the selected $H'$and the parameter s. The probability that $RANDOM$ selects a given price $p = 2^k$ is $\frac{1}{s}\sum_{i=k}^{s+k} \mathbf{Pr}[H' = 2^i] \geq \frac{c}{(s+k)\log^2(s+k)}$. This implies that the approximation ratio, for $s = \log(2mn)$ and $p \leq H$ is $O(\log(nmH)\log^2(\log(mnH)))$.

4 Limited supply setting for bidders with additive valuations

We extend the result in the previous section to the limited-supply setting for the special case of bidders with additive valuations. This is also a setting for which Likhodedov and Sandholm [9] provide logarithmic approximations but via much more complex bundle-pricing mechanisms. We analyze here the mechanism $RANDOM\ Limited\ Supply$ (given below).

Mechanism 2 $RANDOM\ Limited\ Supply$

Input: m = number of bidders, n = number of items, c_j = number of copies of item j available, $H = \max_{i,S}(\mathbf{v}_i(S))$.

Step 1 Let $\tilde{q}_l = \frac{H}{2^{l-1}}$, for $l \geq 1$, $l \in \mathbb{Z}$, and let $s = \log(4nm)$

Step 2 Pick a price $\tilde{q}$ uniformly at random in $\{\tilde{q}_1, ..., \tilde{q}_s\}$, and post the item pricing $\tilde{q}\mathbf{1}_n$.

Step 3 Arrange the bidders in an *arbitrary* order. Allow each in turn to purchase the bundle of highest utility to them from whatever items still remain.

Theorem 3 *The $RANDOM\ Limited\ Supply$ Mechanism is truthful and $O(\log(n) + \log(m))$-competitive with respect to the total social welfare of the bidders.*

Proof: The truthfulness of $RANDOM\ Limited\ Supply$ is immediate. Since all our bidders have additive valuations, we can decompose each bidder's valuation on a set S into a separate valuation $\mathbf{v}_i(\{j\})$ for each item $j \in S$, i.e. $\mathbf{v}_i(S) = \sum_{j \in S} \mathbf{v}_i(\{j\})$, and so the bidder makes an independent decision on each item whether to buy it or not. For any j, let us denote by $W_j = \{w_{j,1}, ..., w_{j,c_j}\}$ the "winners" of item j in the optimal social welfare solution. So, the total social welfare is $\sum_{j=1}^{n}\left[\sum_{i=1}^{c_j} \mathbf{v}_{w_{j,i}}(\{j\})\right]$.

Consider now a single item j, and let c_j be the number of copies of j available. We know that by picking a random price in the range $\{\tilde{q}_1, ..., \tilde{q}_s\}$ our expected profit with respect to the set of winners

$w_{j,1}, ..., w_{j,c_j}$ is at least $\frac{1}{2(2+\log(n)+\log(m))} \sum_{i=1}^{c_j} \mathbf{v}_{w_{j,i}}(\{j\}) - \frac{H}{4n}$. Note that some other bidders not in the set of winners W_j might end up buying item j in our mechanism, but certainly our expected profit is no less than $\frac{1}{2(2+\log(n)+\log(m))} \sum_{i=1}^{c_j} \mathbf{v}_{w_{j,i}}(\{j\}) - \frac{H}{4n}$. So, summing over all items, we obtain that our expected profit is at least $\frac{1}{2(2+\log(n)+\log(m))} \sum_{j=1}^{n} \left[\sum_{i=1}^{c_j} \mathbf{v}_{w_{j,i}}(\{j\})\right] - \frac{H}{4}$, which is at least $\frac{1}{4(2+\log(n)+\log(m))} \sum_{j=1}^{n} \left[\sum_{i=1}^{c_j} \mathbf{v}_{w_{j,i}}(\{j\})\right]$, as desired. ■

5 Lower Bounds

It is interesting to note that we can potentially extract substantially more profit by using bundle pricing rather than item prices. We can show a $\log(n)$ gap between the profit achievable by the best item pricing and the best bundle pricing for the simple case of single minded bidders, even for the case when *all* bidders have valuations 1. Specifically:

Theorem 4 *There exists a* $\log n$ *gap between the profit achievable by best item pricing and the profit achievable by the best bundle pricing, even for the case of single minded bidders when all valuations are* 1.

Proof: Assume $n = 2^r$ items. All bidders will be single-minded with valuation 1 on their desired set. We will have $r + 1$ sets $B_0, B_1, \ldots, B_r$ of n bidders each, so $m = (r + 1)n$ bidders total. The bidders in B_0 all want the entire set $J = \{1, \ldots, n\}$. The bidders in B_1 consist of $n/2$ who want the subset $\{1, \ldots, n/2\}$ and $n/2$ who want the subset $\{n/2 + 1, \ldots, n\}$. The bidders in B_2 consist of $n/4$ who want $\{1, \ldots, n/4\}$, $n/4$ who want $\{n/4+1, \ldots, n/2\}$, $n/4$ who want $\{n/2+1, \ldots, 3n/4\}$, and $n/4$ who want $\{3n/4+1, \ldots, n\}$, and so on. Bidders in B_r each want a distinct single item. Clearly there is a bundle-pricing that achieves revenue $(r + 1)n$, namely just price all bundles at 1.

Next we show that any item pricing has revenue of at most $2n - 1$. We do this by induction on n. The base case ($n = 1$) is trivial. For the general case, first suppose the sum of all prices is at most 1. Then all bidders purchase, but the revenue from bidders in B_0 is at most n, the revenue from bidders in B_1 is at most $n/2$, the revenue from bidders in B_2 is at most $n/4$ and so on, so the total revenue is at most $2n - 1$. On the other hand, suppose the sum of all prices is more than 1. Then we receive no revenue at all from bidders in B_0, and the problem splits into two problems of size $n/2$ of the exact same form. By induction, we can get profit at most $n - 1$ from each, so again we get profit at most $2n - 1$ in total. ■

6 Conclusions

In this paper we exhibit a surprisingly simple mechanism for unlimited-supply combinatorial auctions, which offers the *same* price for each item and achieves revenue within a logarithmic factor of the total social welfare for bidders with *general* valuation functions – which might not even satisfy the free disposal property. This generalizes the prior work of Guruswami et. al [7], who have considered only the special cases of single-minded and unit-demand bidders, as well the results of Likhodedov and Sandholm [9] who have provided logarithmic approximations but via much more complex bundle-pricing mechanisms.

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] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing Truth-telling Online Mechanisms to Online Optimization. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 503–510, 2003.

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

[4] 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.

[5] 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.

[6] 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.

[7] 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.

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

[9] A. Likhodedov and T. Sandholm. Approximating Revenue-Maximizing Combinatorial Auctions. In *The Twentieth National Conference on Artificial Intelligence (AAAI)*, pages 267–274, 2005.

[10] N. Nisan, T. Rougharden, E. Tardos, and V. Vazirani (Eds.). *Algorithmic Game Theory*. 2006. To appear.