Competitive Auctions for Multiple Digital Goods

Andrew V. Goldberg[1] and Jason D. Hartline[2]

[1] STAR Laboratory, InterTrust Technologies Corp.
4750 Patrick Henry Dr., Santa Clara, CA 95054, USA
goldberg@intertrust.com
[2] Computer Science Department, University of Washington.
hartline@cs.washington.edu

Abstract. Competitive auctions encourage consumers to bid their utility values while achieving revenue close to that of fixed pricing with perfect market analysis. These auctions were introduced in [6] in the context of selling an unlimited number of copies of a single item (e.g., rights to watch a movie broadcast). In this paper we study the case of multiple items (e.g., concurrent broadcast of several movies). We show auctions that are competitive for this case. The underlying auction mechanisms are more sophisticated than in the single item case, and require solving an interesting optimization problem. Our results are based on a sampling problem that may have other applications.

1 Introduction

Consider an airplane flight where passengers have individual movie screens and can choose to view one out of a dozen movies that are broadcast simultaneously. The flight is only long enough for one movie to be seen. The airline wants to price movies to maximize its revenue. Currently, airlines charge a flat fee for movies. Even if the fee is based on a careful marketing study, passenger demographics may vary from one flight to another, and individual utilities can vary with flight route, time of the year, etc. Therefore a non-adaptive pricing is unlikely to be optimal for the seller. We investigate adaptive pricing via auctions.

We consider the problem of selling several items, with each item available in unlimited supply. By unlimited supply we mean that either the seller has at least as many items as there are consumers, or that the seller can reproduce items on demand at negligible marginal cost. Of particular interest are digital and broadcast items. With unlimited supply, consumer *utilities*, the maximum price a consumer is willing to pay for an item, are the sole factor determining sale prices and number of items sold. We assume that each consumer has potentially different utilities for different items, and needs one item only. The seller's goal is to set prices to maximize total revenue.

In the scarce supply case, multiple item auctions have been studied by Shapley and Shubik [15]. (See [13] for a survey of the area.) Results for the scarce case, however, do not directly apply to the unlimited supply case. Consider the case where each item for sale is unique – for example the real estate market considered in [15]. In this case consumers will bid heavily for highly desirable items, which will sell for a high price. In contrast, in the unlimited supply case the seller can in principle give every consumer a copy of the item the consumer desires most. However, in such an auction, the consumer has no incentive to bid high. Thus a good auction mechanism must in some cases limit the number of copies of each item.

A consumer's utility value for an item is the most they are willing to pay for that item. We would like to develop auctions in which rational consumers bid their utilities. In game theory, such auctions are called *truthful* and are special cases of strategyproof mechanisms, which have been studied for a long time. For example, the Vickrey–Clarke–Groves mechanism [3, 8, 17] maximizes the general welfare of a system. The Shapley Value [14] mechanism shares costs among the participants. Recent work in the Computer Science community combines economic or game-theoretic questions with computational questions or techniques; see e.g., [5, 10, 9, 12].

Our previous work [6, 7], addressed a special case of the unlimited supply auction problem for a single item. In particular, we introduced *competitive auctions* which are truthful and at the same time attain revenues close to that of fixed pricing with perfect market analysis. As the term suggests, competitive analysis of auctions is similar in spirit to the analysis of on-line algorithms; see, e.g., [1, 16]. We introduced several randomized auctions which are competitive under certain assumptions and showed some impossibility results, including the nonexistence of deterministic competitive auctions.

In this paper we extend some of these results to multiple item auctions. In particular, we develop competitive auctions based on random sampling. These auction mechanisms are intuitive but more sophisticated than in the single item case. We introduce a multiple item variant of the random sampling auction and of the dual price auction and show that these auctions are competitive under certain assumptions. We also discuss a deterministic auction. Although this auction is not competitive in the worst-case, its single item variant worked well in most cases in the experimental study [6, 7]) and show that these auctions are competitive under certain assumptions. We also discuss a deterministic auction. Although this auction is not competitive in the worst-case, its single item variant worked well in most cases in the experimental study [6, 7].

Our work uses the relationship between multiple item auctions and mathematical programming pointed out by Shapley and Shubik. For our random sampling auction we need to solve the following subproblem, which is interesting on its own: given the consumer's utilities, find item prices that maximize seller's revenue. We state this problem as a nonlinear mathematical program.

One of our main results is on a sampling problem that may be of independent interest. A variant of the sampling problem is as follows. Suppose we have n applicants and m tests. Each applicant takes each test and gets a real-valued score. We have to select k applicants based on the results of these scores. Furthermore suppose that we choose a random subset of the applicants, call the applicants in the subset red, and call the remaining applicants blue. After the results of the tests are known and the subset is selected, an adversary selects the k winning applicants while obeying the following restriction: If an applicant x is accepted and for every test, applicant y get a score that is at least as good as the score of x, then y must be accepted as well. Adversary's goal is to bias the admission in favor of red applicants. Although we study a slightly different problem, our techniques can be used to show that if $k = o(m^2 \log n)$, then with high probability the ratio of the number of red applicants to the number of blue applicants is bounded by a constant.

This problem seems natural.One can view candidates as points in m-dimensional space, and view the adversary as selecting a shift of the positive quadrant so that the shifted quadrant contains k points total and as many red points as possible.

In on-line markets, with rapid changes and the availability of computer trading tools and agents, pricing using auctions is sometimes attractive. Competitive auctions for multiple unlimited supply items may be useful in some of these scenarios.

2 Background

The input to an auction is a number of bidders, n, a number of items, m and a set of bids $\{a_{ij}\}$. We assume that all bids are nonnegative and that there is no collusion among the bidders. We study the case when each bidder wants only a single item.

Given a set of bids, the outcome of an auction is an assignment of a subset of (winning) bidders to items. Each bidder i in the subset is assigned a single item j and a sales price of at most a_{ij}. An item can be assigned to any number of bidders. A *deterministic auction mechanism* maps auction inputs to auction outcomes. A *randomized auction mechanism* maps inputs to probability distributions on auction outcomes. We use $\mathcal{R}$ to denote the auction *revenue* for a particular auction mechanism and set of bids. $\mathcal{R}$ is the sum of all sale prices. For randomized auctions, $\mathcal{R}$ is a random variable. We will assume that the m-th item is a dummy item of no value and that all bidders have utility of zero for this item ($a_{im} = 0$ for all i). Losing is then equivalent to being assigned the dummy item at cost zero.

We say that an auction is *single-price* if the sale prices for copies of the same item are the same, and *multiple-price* otherwise.

Next we define *truthful auctions*, first introduced by Vickrey [17]. Let u_{ij} be bidder i's utility value for item j. Define a *bidder's profit* to be the difference between the bidder's utility value for the item won and the price the bidder pays if they win the auction, or zero if they lose. An auction is truthful if bidding u_{ij} is a dominant strategy for bidder i. In other words, the bidder's profit (or expected profit, for randomized auctions), as a function of the bidder's bids $(a_{i1}, \dots, a_{im})$, is maximized at the bidder's utility values $(u_{i1}, \dots, u_{im})$, for any fixed values of the other bidders' bids. Truthfulness is a strong condition for auctions: bidding utility maximizes the profit of the bidder no matter what the other bidders' strategies are. When considering truthful auctions, we assume that $a_{ij} = u_{ij}$, unless mentioned otherwise.

To enable analysis of auction revenue we define several parameters of an input set of bids. The revenue for optimal fixed pricing is $\mathcal{F}$. Note that $\mathcal{F}$ can also be interpreted as the revenue due to the optimal nontruthful single-price auction. Other parameters that we use in analysis are ℓ, the lowest bid value, and h, the highest bid value. Because bids can be arbitrarily scaled, we assume, without loss of generality, that $\ell = 1$, in which case h is really the ratio of the highest bid to the lowest bid.

Analogous to on-line algorithm theory, we express auction performance relative to that for the optimal nontruthful auction, as ratios $\mathcal{R}/\mathcal{F}$. However, we solve a maximization problem, while on-line algorithms solve minimization problems. Thus, positive results, which are lower bounds on $\mathcal{R}/\mathcal{F}$, are expressed using "Ω".

Note that h and $\mathcal{F}$ are used only for analysis. Our auctions work without knowing their values in advance.

As shown in [6], if we do not impose any restrictions on h, we get the upper bound of $\mathcal{R}/\mathcal{F} = O(1/h)$. To prevent this upper bound on auction revenue we can make the assumption that the optimal revenue $\mathcal{F}$ is significantly larger than h, the highest bid. With this assumption, optimal fixed pricing sells many items.

We say that an auction is *competitive* under certain assumptions if when the assumptions hold, the revenue is $\Omega(\mathcal{F})$.

For convenience, we assume that the input bids are *non-degenerate*, i.e., all input bids values a_{ij} are distinct or zero. This assumption can be made without loss of generality because we can always apply a random perturbation or use lexicographic tie-breaking to achieve it.

As shown for the single-commodity case [6], no deterministic auction is competitive in the worst case. Our competitive auctions are randomized. We use the following lemma, which is a variation of the Chernoff bound (see e.g. [2, 11]), as the main tool in our analysis.

Lemma 1. *Consider a set A and its subset $B \subset A$. Suppose we pick an integer k such that $0 < k < |A|$ and a random subset (sample) $S \subset A$ of size k. Then for $0 < \delta \leq 1$ we have*

$$\mathbf{Pr}[|S \cap B| < (1 - \delta)|B| \cdot k/|A|] < \exp(-|B| \cdot k\delta^2/(2|A|)).$$

Proof. We refer to elements of A as points. Note that $|S \cap B|$ is the number of sample points in B, and its expected value is $|B| \cdot k/|A|$. Let $p = k/|A|$. If instead of selecting a sample of size exactly k we choose each point to be in the sample independently with probability p then the Chernoff bound would yield the lemma.

Let $A = \{a_1, \ldots, a_n\}$ and without loss of generality assume that $B = \{a_1, \ldots, a_k\}$. We can view the process of selecting S as follows. Consider the elements of A in the order induced by the indices. For each element a_i considered, select the element with probability p_i, where p_i depends on the selections made up to this point.

At the point when a_{i+1} is considered, let t be the number currently selected points. Then $i - t$ is the number of points considered but not selected. Suppose that $t/i < p$. Then $p_{i+1} > p$.

We conclude that when we select the sample as a random subset of size k, the probability that the number of sample points in B is less than the expected value is smaller than in the case we select each point to be in the sample with probability p. ■

3 Fixed Price Auction and Optimal Prices

Consider the following *fixed price auction*. The bidders supply the bids and the seller supplies the *sale prices*, r_j, $1 \leq j \leq m$. Define $c_{ij} = a_{ij} - r_j$. The auction assigns each bidder i to the item j with the maximum c_{ij}, if the maximum is nonnegative, and to no item otherwise. In case of a tie, we chose the item with the maximum j. If a bidder i is assigned item j, the corresponding sale price is r_j.

Lemma 2. *Suppose the sale prices are set independently of the input bids. Then the fixed price auction is truthful.*

Proof. If bidder i gets object j, the bidder's price is at least r_j and the bidder's profit is at most $a_{ij} - r_j$. The best possible profit for i is $\max_j(u_{ij} - r_j)$. If the bidder bids $a_{ij} = u_{ij}$, this is exactly the profit of the bidder. ■

Remark Although we assume that the bidders do not see sale prices before making their bids, the lemma holds even if the bidders do see the prices.

Now consider the following *optimal pricing* problem: Given a set of bids, find the set of prices such that the fixed price auction brings the highest revenue. Suppose an auction solves this problem and uses the resulting prices. We call this auction the *optimal nontruthful single-price auction* and denote its revenue by $\mathcal{F}$. We can interpret $\mathcal{F}$ as the revenue of fixed pricing using perfect market analysis or as the revenue of the optimal nontruthful single-price auction. The prices depend on the input bids, and one can easily show this auction is nontruthful.

We use $\mathcal{F}$ to measure performance of our truthful auctions. Although one might think that being a single-price auction is a serious restriction, in the single-item auction case this is not so. In this case, the revenue of the optimal single-price auction is at least as big as the expected revenue of any reasonable[1] (possible multiple-price) truthful auction; see [6].

Next we state the optimal pricing problem as a mathematical programming problem. We start by stating the problem of finding a bidder-optimal object assignment given the bids and the sale prices as an integer programming problem. This problem is a special case of the b-matching problem [4] (bipartite, weighted,

[1] See [6] for the precise definition of reasonable. The intuition is that we preclude auctions that are taylored to specific inputs. Such an auction would perform well these specific inputs, but poorly on all others.

and capacitated, with unit node capacities on one side and infinite capacities on the other). For the limited supply case, when only one copy of an item is available, the classical paper [15] takes a similar approach. For our case, this problem is easy to solve by taking the maximum as in the previous section. However, we then treat sale prices as variables to get a mathematical programming formulation of the optimal pricing problem.

One can show that the optimal price problem is equivalent to the following mathematical programming problem; we omit details.

$$\begin{array}{lrl} \max & \sum_j \sum_i x_{ij} r_j & \text{subject to} \\ & r_m = 0 & \\ & \sum_j x_{ij} \leq 1 & 1 \leq i \leq n \\ & x_{ij} \geq 0 & 1 \leq i \leq n,\ m \leq j \leq m \\ & p_i + r_j \geq a_{ij} & 1 \leq i \leq n,\ m \leq j \leq m \\ & \sum_i p_i = \sum_j \sum_i x_{ij} \cdot (a_{ij} - r_j) & \end{array} \tag{1}$$

This problem has quadratic objective function; some constraints are linear while other constraints are quadratic. Here x_{ij} is one exactly when bidder i gets item j and p_i's are profits of the corresponding bidders.

Since $\sum_j \sum_i x_{ij} r_j = \sum_j \sum_i x_{ij} a_{ij} - \sum_i p_i \leq \sum_j \sum_i a_{ij}$, the objective function is bounded. Since the feasibility region is closed, it follows that (1) always has an optimal solution.

We omit proofs of the next two results.

Lemma 3. *For any solution of (1) with fractional x_{ij}'s there is a solution with $x_{ij} \in \{0, 1\}$ and an objective function value that is at least as good.*

Theorem 1. *Consider sale prices defined by an optimal solution of (1). The revenue of the fixed price auction that uses these prices and has bids a_{ij} in the input is equal to the objective function value of the optimal solution.*

Recall that we use the problem (1) to find a set of prices that maximizes the fixed price auction revenue. In the rest of the paper we assume that we can compute such prices and leave open the question of how to do this efficiently. Note that we could also use an approximate solution.

4 The Random Sampling Auction

We use random sampling to make the optimal single-price auction truthful.

The *random sampling auction* works as follows.

1. Pick a random sample S of the set of bidders. Let N be the set of bidders not in the sample.
2. Compute the optimal sale prices for S as outlined in the previous section.
3. The result of the random sampling auction is then just the result of running the fixed-price auction on N using the sale prices computed in the previous step. All bidders in S lose the auction.

The sample size is a tradeoff between how well the sample represents the input and how much potential revenue is wasted because all bidders in the sample lose. Unless mentioned otherwise, we assume that the sample size is $n/2$ or, if n is odd, the floor or the ceiling of $n/2$ with probability $1/2$.

The facts that the bidders who determine the prices lose the auction and that the fixed price auction is truthful imply the following result.

Lemma 4. *The random sampling auction is truthful.*

Remark Another natural way of sampling is to sample bids instead of bidders. However, this does not seem to lead to a truthful auction, because bidder's bids selected in the sample may influence the price used to satisfy the bidder's remaining bids.

Next we show that, under certain assumptions, the auction's revenue $\mathcal{R}$ is within a constant factor of $\mathcal{F}$. Without loss of generality, for every $1 \le i \le n, 1 \le j \le m$, if a_{ij} is undefined (not in the input) we define a_{ij} to be zero. For every bidder i, we view $(a_{i1}, \ldots, a_{im})$ as a point in the m-dimensional space and denote this point by v_i. Thus v_i is in the quadrant Q of the m-dimensional space where all coordinates are nonnegative. We denote the set of all input points by B.

For a fixed m and a set of sale prices $r_1, \ldots, r_m$, let R_j be a region in the m-dimensional space such that if $v_i \in R_j$, then i prefers j to any other item, i.e., for any $1 \le k \le m$, $c_{ij} \ge c_{ik}$ (recall that $c_{ij} = a_{ij} - r_j$). We would like $\{R_j : 1 \le j \le m\}$ to be a partitioning of Q. We achieve this by assigning every boundary point to the highest-index region containing the point. (This is consistent with our tie-breaking rule for the fixed price auction.) R_j is a convex (and therefore connected) region in Q. In fact, the region R_j is as follows:

$$R_j = \{x : x_j \ge r_j \ \& \ x_j - r_j \ge x_k - r_k \ \forall k \ne j\}. \tag{2}$$

Figure 1 shows a two item auction with prices r_1 and r_2 for items 1 and 2 respectively. These prices induce the regions $R_1 = R_1' \cup R_1''$ and $R_2 = R_2' \cup R_2''$. Arrows point to selling prices for the bidders in each region.

Thus sampling and computing r_j's partitions Q into the regions, and each bidder i in N gets the item corresponding to the region that i is in. Intuitively, our analysis says that if a region has many sample points, it must have a comparable number of nonsample points – even though the regions are defined based on the sample. The latter fact makes the analysis difficult by introducing conditioning. Intuitively, we deal with the conditioning by considering regions defined by the input independently of the sample.

For a given input, let $q_1, \ldots, q_m$ be a set of optimal prices for the input bids that yield revenue $\mathcal{F}$. These prices induce the regions discussed above. Bidders in region R_j pay q_j for the item j. If we sample half of the points, the expected number of sample points in a region R_j is half of the total number of points in the region, and for the prices $q_1, \ldots, q_m$, the expected revenue is $\mathcal{F}/2$. The optimal fixed pricing on the sample does at least as well. Thus the expected revenue of optimal fixed pricing of the sample, $\mathbf{E}[\mathcal{F}_s]$, is at least $\mathcal{F}/2$. However,

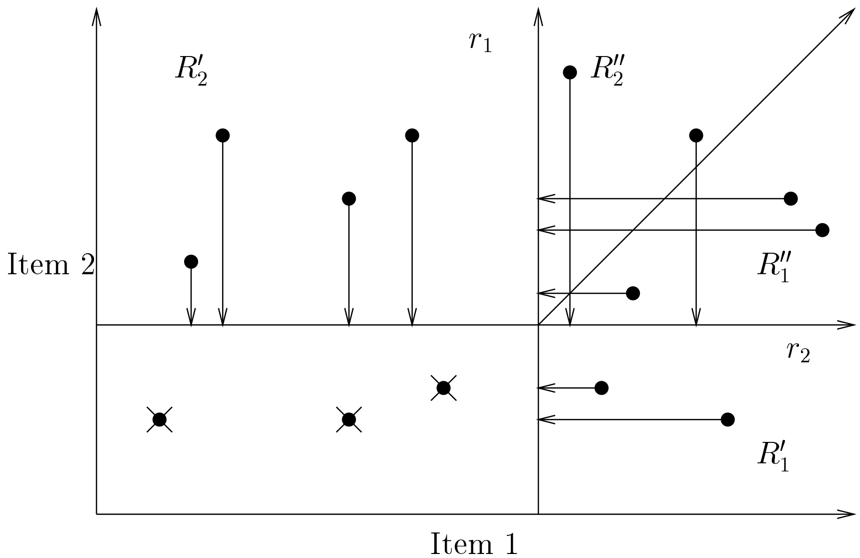

Fig. 1. Two item auction with regions R_1 and R_2

we need a high-probability result. Our goal is to show that with high probability $\mathbf{E}[\mathcal{F}_s]$ is close to $\mathcal{F}/2$ and that $\mathbf{E}[\mathcal{R}]$ is close to $\mathbf{E}[\mathcal{F}_s]$, where $\mathcal{R}$ is the revenue of the random sampling auction.

We say that a set $A \subseteq B$ is *t-feasible* if A is nonempty and for some set of sale prices, A is exactly the set of points in R_t. For each feasible set A, we define its signature $S_A = (s_1, \ldots, s_m)$ such that s_i's are (not necessarily distinct) elements of A and, for a fixed t, different t-feasible sets have different signatures. In the following discussion, s_{ij} denotes the j-th coordinate of s_i.

We construct signatures as follows. Let R_t be a region defining A. R_t is determined by a set of prices $(r_1, \ldots, r_m)$. We first increase all r_j's by the same amount (moving R_t diagonally) until some point in A is on the boundary of R_t. Note that since we change all prices by the same amount, the limiting constraint from (2) is $x_t \geq r_t$. Thus the stopping is defined by $x_t = r_t$, and the point on the boundary has the smallest t-th coordinate among the points in A. We set s_t to this point.

Then for $j \neq t$, we move the the region starting at its current position down the j-th coordinate direction by reducing r_j until the first point hits the boundary. The boundary we hit is defined by $x_t - r_t = x_j - r_j$, and the point that hits it first has the minimum $x_j - x_t + s_{tt}$ among the points in A. Observe that the point s_t remains on the boundary $x_t = r_t$, and therefore we stop before r_j becomes negative. When we stop, we take a point that hits the boundary and assign s_j to it.

Consider the set of points in the signature, $S_A = \{s_1, \ldots, s_m\}$. Define R to be the region we got at the end of the procedure that computed S_A. R is defined by

$$R = \{x : x_t \geq s_{tt} \ \& \ x_t - s_{tt} \geq x_j - s_{jj} \ \forall j \neq t\}.$$

It follows that R can be constructed directly from S_A.

Figure 2 shows the signatures we get from the prices r_1 and r_2. The points on the boundary of the shaded region are the signature of that region. Note, for example, that there are no points in R_1 that are not inside the boundary induced by the signature for R_1.

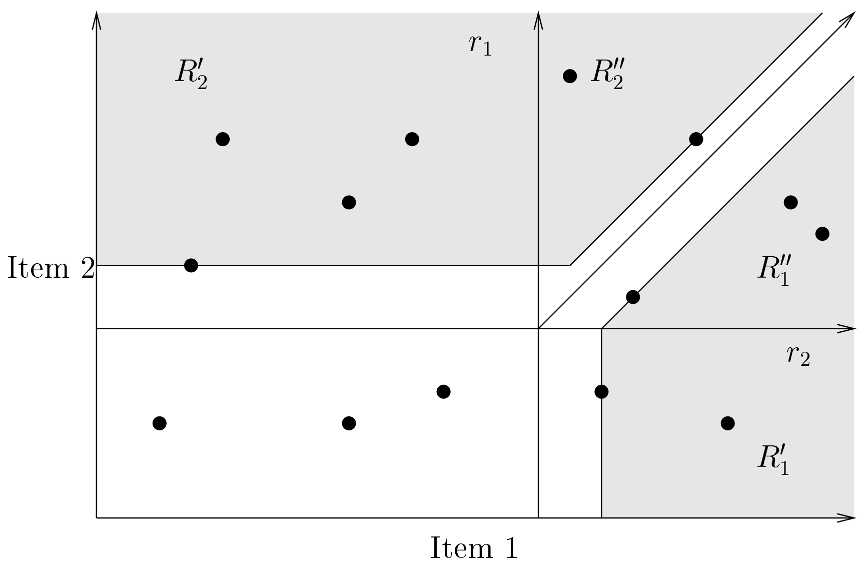

Fig. 2. Signatures in a two item auction

Suppose two feasible sets have the same signature S and let R be the region defined by the signature. Then the two sets are exactly the set of points in R, and are thus identical.

The next two lemmas are simple, so we omit the proofs.

Lemma 5. *For each t, $1 \leq t \leq m$, there are at most n^m t-feasible sets.*

Lemma 6. *For every t-feasible subset C of the sample S there is a t-feasible subset A of the input such that $C = A \cap S$.*

For $k \geq 1$ and $0 < \delta < 1$, we say that a sample S is (k, δ)-*balanced* if for every $1 \leq t \leq m$ and for every t-feasible subset of the input, A, such that $|A| \geq k$, we have

$$(1 - \delta) \leq (|A \cap S|)/(|A \cap N|) \leq 1/(1 - \delta).$$

Lemma 7. *The probability that a sample containing half of the input points is (k, δ)-balanced is at least $1 - 2mn^m \exp(-k\delta^2/8)$.*

Proof. Lemma 1 implies that the probability that for a set A with $|A| \geq k$,

$$\mathbf{Pr}[|A \cap S| < (1 - \delta)|A \cap N|] < \exp(-k\delta^2/8)$$

and

$$\mathbf{Pr}[|A \cap N| < (1 - \delta)|A \cap S|] < \exp(-k\delta^2/8).$$

Note that the fact that the number of sample points in one subset is close to its expectation makes it no less likely that the number of sample points in another subset is close to expectation. Thus the conditioning we get is favorable. By Lemma 6, there are at most n^m t-feasible subsets for every t, so the total number of feasible subsets is mn^m. These observations imply the lemma. ∎

Theorem 2. *Assume $\alpha h m^2 \ln n \leq \mathcal{F}$ and $m \geq 2$. Then $\mathcal{R} \geq \mathcal{F}/24$ with probability of at least $1 - \exp(-\alpha/1728)$ (for some constant $\alpha > 1$).*

Proof. Consider Lemma 7 with $\delta = 1/2$ and $k = \alpha m \log n/12$. The probability that the sample is (k, δ)-balanced is

$$1 - 2mn^m \exp(-k\delta^2/8) = 1 - 2mn^m \exp(-\alpha m \log n/864) \geq 1 - \exp(-\alpha/1728)$$

for $m \geq 2$. For the rest of the proof we assume that the sample is (k, δ)-balanced; we call this the balanced sample assumption.

Next we show that the revenue of the auction on the sample, $\mathcal{F}_s$, satisfies $\mathcal{F}_s \geq \mathcal{F}/6$. Let Q_i be the set of bidders who get item i when computing $\mathcal{F}$ on the entire bid set. Consider sets Q_i containing less than $(\alpha m \log n)/2$ bidders. The total contribution of such sets to $\mathcal{F}$ is less then $\mathcal{F}/2$. This is because there are at most m such sets and each bid is at most h giving a maximum possible revenue of $\alpha h m^2 \log n/2 = \mathcal{F}/2$. Thus the contribution of the sets with at least $(\alpha m \log n)/2$ bidders is more than $\mathcal{F}/2$, and we restrict our attention to such sets. By the balanced sample assumption, each such set contains at least $1/3$ sample points, and thus $\mathcal{F}_s \geq (1/3)\mathcal{F}/2 = \mathcal{F}/6$.

Finally we show that $\mathcal{R} \geq \mathcal{F}/24$ using a similar argument. Let R_i be the regions defined by the prices computed by the auction on the sample. Consider the regions containing less than $(\alpha m \log n)/12$ sample points. The total contribution of such sets to the revenue is less then $\mathcal{F}/12$. The remaining regions contribute at least $\mathcal{F}/12$ (out of $\mathcal{F}/6$). Each remaining region contains at least $(\alpha m \log n)/12$ sample points. By the balanced sample assumption, each such region contains at least one nonsample point for every two sample point, and thus $\mathcal{R} \geq \mathcal{F}/24$. ∎

Lemma 4 and Theorem 2 imply that if the assumptions of the theorem hold, the random sampling auction is competitive.

4.1 The Dual Price Auction

The random sampling auction is wasteful in the sense that all bidders in the sample lose the auction. The *dual price auction* eliminates the waste by treating

S and N symmetrically: S is used to compute sale prices for N and vice versa. Note that for each item, the two sale prices used are, in general, different; this motivates the name of the auction.

By symmetry, the expected revenue of the dual price auction is twice the expected revenue of the single price auction with $n/2$ sample size. Thus, under conditions of Theorem 2 the dual price auction is competitive.

5 A Deterministic Auction

The following auction is a generalization of the *deterministic optimal threshold auction* introduced in [6] to the multi-item case. Although not competitive in general, the single-item variant of this auction works well when the input is non-pathological, e.g., when bidder utilities are selected independently from the same distribution.

The deterministic auction determines what item, if any, the bidder i gets as follows. It deletes i from B, computes optimal prices for the remaining bidders, and then chooses the most profitable item for i under these prices. This is done independently for each bidder. This auction is truthful but, as we have mentioned, not competitive in some cases.

6 Concluding Remarks

Our analysis of the random sampling auction is somewhat brute-force, and a more careful analysis may lead to better results, both in terms of constants and in terms of asymptotic bounds. In particular, the assumption $\alpha h m^2 \ln n \leq \mathcal{F}$ in Theorem 2 may be stronger than necessary. One can prove that $\mathcal{F}_s = \Omega(\mathcal{F})$ assuming $\alpha h m \leq \mathcal{F}$. We wonder if the theorem holds under this weaker assumption.

Although our theoretical bounds require m to be small compared to n and the optimal fixed price solution to contain a large number of items, it is likely that in practice our auctions will work well for moderately large m and moderately small optimal fixed price solutions. This is because our analysis is for the worst-case. In many real-life applications, bidder utilities for the same item are closely correlated and our auctions perform better.

The optimal fixed pricing problem has a very special form that may allow one to solve this problem efficiently. Note that if one uses an approximation algorithm to solve the problem (say within 2% of the optimal) and our auctions remain truthful. (This is in contrast to combinatorial auctions [9].) It is possible that in practice this problem can be solved approximately, in reasonable time, using general nonlinear optimization techniques. We leave an existence of such an algorithm as an open problem.

Another open problem is a generalization of our results. One possible generalization is to the case when some items are in fixed supply. Another generalization is to the case when consumer i wants up to k_i items.

References

1. A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis.* Cambridge University Press, 1998.
2. H. Chernoff. A Measure of Asymptotic Efficiency for Test of a Hypothesis Based on the Sum of Observations. *Anals of Math. Stat.*, 23:493–509, 1952.
3. E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.
4. J. Edmonds and E. L. Johnson. Matching, a Well-Solved Class of Integer Linear Programs. In R. Guy, H. Haneni, and J. Schönhein, editors, *Combinatorial Structures and Their Applications*, pages 89–92. Gordon and Breach, NY, 1970.
5. J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. In *Proc. of 32nd Symposium Theory of Computing*, pages 218–226. ACM Press, New York, 2000.
6. A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. Technical Report STAR-TR-99.09.01, STAR Laboratory, InterTrust Tech. Corp., Santa Clara, CA, 1999. Available at URL http://www.star-lab.com/tr/tr-99-01.html.
7. A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Alg.*, pages 735–744. ACM/SIAM, 2001.
8. T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.
9. D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. of 1st ACM Conf. on E-Commerce*, pages 96–102. ACM Press, New York, 1999.
10. N. Linial. Game Theoretic Aspects of Computing. In *Handbook of Game Theory*, volume 2, pages 1339–1395. Elseveir Science Publishers B.V., 1994.
11. R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, 1995.
12. N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proc. of 31st Symposium on Theory of Computing*, pages 129–140. ACM Press, New York, 1999.
13. A. E. Roth and M. A. Oliveira Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis.* Cambridge University Press, 1990.
14. L. S. Shapley. Core of Convex Games. *Int. J. of Game Theory*, 1:11–26, 1971.
15. L. S. Shapley and M. Shubik. The Assignment Game I: The Core. *Int. J. of Game Theory*, 1:111–130, 1972.
16. D. D. Sleator and R. E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Communications of the ACM*, 28:202–208, 1985.
17. W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.