Competitive Auctions and Digital Goods

Andrew V. Goldberg[*] Jason D. Hartline[†] Andrew Wright[*]
goldberg@intertrust.com hartline@cs.washington.edu wright@intertrust.com

Abstract

We study a class of single round, sealed bid auctions for items in unlimited supply such as digital goods. We focus on auctions that are truthful and competitive. Truthful auctions encourage bidders to bid their utility; competitive auctions yield revenue within a constant factor of the revenue for optimal fixed pricing. We show that for any truthful auction, even a multi-price auction, the expected revenue does not exceed that for optimal fixed pricing. We also give a bound on how far the revenue for optimal fixed pricing can be from the total market utility. We show that several randomized auctions are truthful and competitive under certain assumptions, and that no truthful deterministic auction is competitive. We present simulation results which confirm that our auctions compare favorably to fixed pricing. Some of our results extend to bounded supply markets, for which we also get truthful and competitive auctions.

1 Introduction.

Consider the problem of selling a number of identical items to consumers who each want a single item and the items are 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 items such as downloadable audio files and pay-per-view movies. With unlimited supply, consumer *utilities*, the maximum amounts that consumers are willing to pay for the item, are the sole factor determining sale prices. The seller's goal is to maximize their total revenue.

One way to set prices for items in unlimited supply is to estimate consumer utility via market analysis and then set a fixed price. We refer to this method as *fixed pricing*. Pay-per-view movies are an example of fixed pricing for an unlimited supply market. With perfect knowledge of consumer utilities, *optimal fixed pricing* maximizes fixed-price revenue by selecting the optimal price at which to sell items. Fixed pricing generally cannot achieve this ideal due to the inherent inaccuracy of market analysis. If the price is set too high, not enough items may be sold; if the price is set too low, insufficient revenue may be collected per item.

Auctions automatically adjust prices to market conditions. We study single round, sealed bid auctions. Such auctions have been studied for items available in *scarce supply*, where maximizing the revenue requires that all available items be sold. They are especially practical when the number of consumers is large. In particular, Vickrey [17] introduced a multi-item auction that is *truthful*. A truthful auction encourages bidders to bid their utility value. In an untruthful auction, bidders may bid significantly below their utility values, reducing auction revenue.

In a truthful auction, rational bidders bid their utilities. In addition, we would like such an auction to be *competitive*: it must yield revenue within a constant factor of optimal fixed pricing. To be competitive, a truthful auction must vary how many items are sold depending on the bid values. For example, as we show in Section 3, the multi-item Vickrey auction is not competitive if the seller chooses the number of items to sell before knowing the bid values (and not truthful if the seller chooses the number of items after knowing the bid values). As with fixed pricing, selling too few or too many items may not maximize revenue. Thus, the method for choosing how many bids to satisfy is an integral part of a truthful competitive auction.

To our knowledge, auctions have never been studied in a competitive framework. Nor are any existing auctions competitive in the sense that we introduce here. As we explore in Sections 11 and 12, this competitive framework is useful in studying any kind of auction where identical goods are being sold, not just auctions for unlimited supply.

Auctions are becoming a popular pricing mechanism in electronic commerce, both for human users and for trading agents (bots). In many cases, the use of auctions is complicated by the fact that a good bidding strategy for a buyer requires an understanding of strategies and utilities of other buyers. Truthful auction mechanisms may be attractive in this context

[*]STAR Lab., InterTrust Technologies Corp., 4750 Patrick Henry Dr., Santa Clara, CA 95054.

[†]Computer Science Department, University of Washington. Work done while the author was employed by InterTrust.

because they avoid this complication.

In this paper we study a class of truthful auctions for unlimited supply. We study both *single-price* auctions, where every winning bidder pays the same price, and *multi-price* auctions, where the prices may differ. In addition to *deterministic* auctions, we study *randomized* auctions that use randomization to decide which bids to fill and at what price. We develop techniques for design and analysis of auctions for unlimited supply. Our approach is reminiscent of competitive analysis of on-line algorithms [1, 15], where performance of an on-line algorithm is gauged in terms of performance of an optimal off-line algorithm. Here the optimal off-line algorithm is analogous to the optimal fixed pricing mechanism.

Although we develop our results for unlimited supply, some of the results extend to *bounded supply*, where the number of items for sale is bounded, but maximizing revenue might not result in all items being sold. We discuss bounded supply in Section 11.

We view auctions as algorithms for deciding which input bids to fill, and at what price. As with any algorithm, one needs to address the issues of correctness, efficiency, and performance. In the context of this paper, an auction is correct if the auction is truthful and fills each winning bid at or below the bid value. Efficiency of an auction refers to the time needed to process bids. The auctions introduced in this paper are very fast; sorting of the input bids is the most expensive computational operation we perform. As discussed above, we measure auction performance by its revenue relative to the fixed pricing revenue. This computer science approach allows us not only to design new truthful auctions, but also to give theoretical guarantees for their performance. Such provable performance guarantees are new to the area of auction mechanism design.

To state our results more formally, we introduce the following notation. Let n denote the number of bidders. Without loss of generality, we assume that the lowest bid is one and denote the highest bid by h. Let the total utility $\mathcal{T}$ be the sum of all bidders' utilities. $\mathcal{T}$ is an obvious upper bound on the revenue that can be obtained from this set of bidders. Let $\mathcal{F}$ be the optimal fixed pricing revenue. Clearly $\mathcal{F} \leq \mathcal{T}$. $\mathcal{F}$ is an upper bound on the revenue that can be obtained by any fixed priced sale or any single-price auction. We want revenues of our truthful auctions to be competitive with $\mathcal{F}$. We assume that h is small compared to $\mathcal{F}$.[1] This assumption prevents a trivial upper bound on the revenue; see Section 2.

We state some of our results in terms of the total utility, $\mathcal{T}$, and others in terms of the optimal fixed pricing revenue, $\mathcal{F}$. We show that $\mathcal{F}$ compares favorably to $\mathcal{T}$; specifically, $\mathcal{F} = \Omega(\mathcal{T}/\log h)$, and also $\mathcal{F} = \Omega(\mathcal{T}/\log n)$. This result shows that the optimal fixed pricing revenue is within a $\min(\log h, \log n)$ factor of the revenue of any pricing scheme. We use this result to relate various bounds expressed in terms of $\mathcal{T}$ to those expressed in terms of $\mathcal{F}$.

We introduce a class of truthful single-price auctions and a class of truthful multi-price auctions. A randomized auction from the first class is competitive: it has an expected revenue of $\Omega(\mathcal{F}) = \Omega(\mathcal{T}/\log h)$. A dual-price variant of this auction has revenue that is close to $\mathcal{F}$ if $\mathcal{F}/h$ is large enough. A randomized multi-price auction from the second class has an expected revenue of $\Omega(\mathcal{T}/\log h)$. Thus, both of these auctions have the same worst-case bound in terms of $\mathcal{T}$. However, we show that the latter auction is not competitive. Its expected revenue is $\Omega(\mathcal{F}/\sqrt{\log h})$, and this bound is tight: on certain inputs, the expected revenue is $O(\mathcal{F}/\sqrt{\log h})$.

This provides support for using $\mathcal{F}$, rather than $\mathcal{T}$, to define competitive auctions. Another result provides further support. We show that for any truthful auction, even a multi-price one, the expected revenue does not exceed $\mathcal{F}$. This result is somewhat surprising: using single-price auctions does not hurt revenue by more that a constant factor.

A natural question to ask is if there is a truthful deterministic auction with an $\Omega(\mathcal{F})$ revenue. We show that there is none: for any such auction, there is a set of bids that leads to an $O(\mathcal{F}/h)$ revenue, i.e., revenue that is a small fraction of $\mathcal{F}$ if h is large. Thus, for worst-case performance, randomized auctions yield better revenue than the deterministic ones.

Our theoretical analyses are limited in that their performance metrics are accurate only up to a constant factor. As a result, the analyses do not reveal whether one of our auctions dominates the others, or which auction is better for a natural distribution. As a supplement to our theoretical results, we performed a number of simulations to compare our auctions with each other and to fixed pricing on a variety of input families. Our simulations suggest that, on natural inputs, some of our auctions attain revenue very close to $\mathcal{F}$ if the number of bids is large enough. Furthermore, our auctions can outperform fixed pricing with market analysis unless that analysis is fairly accurate. We also show a deterministic auction that, despite the worst-case result, does very well on natural inputs.

We develop a framework for a theoretical and experimental analysis of revenue-maximizing truthful auctions and introduce auctions that perform well in

[1] Some of our results hold under weaker assumptions.

this framework. We show how algorithm analysis techniques can be used within this framework to obtain results that are interesting and in some cases surprising.

Due to the space limit, we omit some material from this version of the paper. The full paper appears at `http://www.intertrust.com/tr/tr-99-01.html`.

2 Competitive Analysis.

We consider auctions with n bidders, each bidder i having a utility value u_i and bidding b_i. We also assume that the bids are ordered so that $b_i \leq b_{i+1}$. In auctions where ties need to be broken, we can assume an arbitrary total order on the bid values that respects the partial order. That is, we can assume that the order given by the indices is strict. We assume that there is no collusion among the bidders.

Given a set of bids, the outcome of an auction is the subset of bids that are satisfied and a corresponding set of sale prices such that, for each winning bid b_i, the associated sale price is at most b_i. A *deterministic auction mechanism* maps sets of bids to auction outcomes. A *randomized auction mechanism* maps sets of bids 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 say that an auction is *truthful* if bidding u_i is a dominant strategy for bidder i. More specifically, let a *bidder's profit* be the difference between the bidder's utility value and the price the bidder pays if they win the auction, or zero if they lose. Then an auction is truthful if a bidder's profit (or expected profit, for randomized auctions), as a function of the bidder's bid, is maximized at the bidder's utility value, for any fixed values of the other bidders' bids. Truthful auctions encourage bidding at utility value if the bid value that maximizes the (expected) profit is unique. When considering truthful auctions, we assume that $b_i = u_i$.

To enable analysis of auction revenue we define several properties of an input set of bids. As stated in the introduction, $\mathcal{T}$ is the sum of all of the bids. An equivalent definition of $\mathcal{T}$ is the revenue due to the optimal multi-price untruthful auction, the one that satisfies all bids at their value. The revenue for optimal fixed pricing is $\mathcal{F}$. Note that $\mathcal{F}$ can also be interpreted as the revenue due to the optimal single-price untruthful auction. More discussion of $\mathcal{F}$ and its relation to the optimal single-price untruthful auction is given in Section 4. The other property of the bids that we use is h, the highest bid value. We assume that bids are scaled such that the lowest bid value is 1.

Analogous to on-line algorithm theory, we express auction performance relative to that for optimal untruthful auctions, as ratios $\mathcal{R}/\mathcal{T}$ or $\mathcal{R}/\mathcal{F}$. However, we solve a maximization problem while on-line algorithms solve minimization problems. Thus, positive results, which are lower bounds on $\mathcal{R}/\mathcal{T}$ or $\mathcal{R}/\mathcal{F}$, are expressed using "Ω". Impossibility results, which are upper bounds on $\mathcal{R}/\mathcal{T}$ or $\mathcal{R}/\mathcal{F}$ for any auction in a certain class, are expressed using "O".

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

If we do not impose any restrictions on h, we get the upper bound of $\mathcal{R}/\mathcal{T} = O(1/h)$. To see this, imagine $n-1$ bids at value 1 and one bid, b_n, at h. An auction that wishes to do better than $O(1/h)$ must base the selling price on bidder n's bid. However, this would encourage bidder n to bid below u_n.

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. That is, for some constant α, we assume that $\alpha h \leq \mathcal{F}$. With this assumption, optimal fixed pricing sells at least α items. In some cases α is a fixed constant. In other cases, success probability approaches 1 as $\alpha \to \infty$.

We say that an auction is *competitive* under certain assumptions if, when the assumptions hold, the auction's revenue is $\Omega(\mathcal{F})$, or equivalently $\mathcal{R}/\mathcal{F} = \Omega(1)$.

3 Prior Results.

Our results are related to the field of mechanism design that combines microeconomic motivation with game-theoretic tools and includes auction mechanisms. For introduction to the area, see for example [10, 13]. In particular, auctions for scarce supply markets have been extensively studied. See [14] for a survey. Some work in the Computer Science community combines economic or game-theoretic questions with computational questions. Earlier results are surveyed in [9]; for more recent results, see e.g. [4, 8, 12].

The k-item Vickrey auction [17] was a starting place for our work. The k-item Vickrey auction is a single-price auction that sells k items to the k highest bidders at the price equal to the $k+1$ highest bid (b_{n-k-1}). For generalizations of Vickrey auctions to the multiple resource case, see e.g. [3, 6, 16].

In the unlimited supply case, taking $k = f(n)$ yields a truthful auction for any function f with $1 \leq f(n) < n$. This auction mechanism is not competitive, however. To see this, consider a *bipolar* input that has k bids at h and $n-k$ bids at 1. In this case $\mathcal{R} = k$ and $\mathcal{F} \geq hk$. This gives an $\mathcal{R}/\mathcal{F} = O(1/h)$ bound. As we show later no deterministic auction can do much better on worst-case distributions, but randomized auctions can.

4 Optimal Untruthful Auctions.

In this section we study the two untruthful auctions, the optimal multi-price auction and the optimal single-price auction, and establish the relationship between their revenues, $\mathcal{T}$ and $\mathcal{F}$. We show that $\mathcal{F} \geq \mathcal{T}/(2\log h)$. This bounds the penalty for requiring auctions to be single-price and allows us to compare bounds expressed in terms of $\mathcal{T}$ with those expressed in terms of $\mathcal{F}$.

To get a better understanding of how the optimal single-price auction works, we define the *optimal threshold function*, $\mathrm{opt}(B)$. This function on a set of bids B returns the fixed price at which items should be sold to achieve revenue $\mathcal{F}$. In the optimal single-price auction, all bid values that are at least $\mathrm{opt}(B)$ will be satisfied at price $\mathrm{opt}(B)$. All other bids will be rejected. Formally,

$$\mathrm{opt}(B) = \mathrm{argmax}_{b_i \in B}\, b_i \cdot (n - i + 1).$$

Note that $n - i + 1$ is the number of bids that are at least b_i. The main result of this section is as follows.

THEOREM 4.1. $\mathcal{F} \geq \mathcal{T}/(2\log h)$.

Proof. Divide the bids into $\log h$ bins by partitioning the bids at values of powers of two. Thus, each bid is less than a factor of two from any other bid in the same bin. Since the sum of the bids is $\mathcal{T}$ and there are $\log h$ bins, then some bin has a sum of at least $\mathcal{T}/\log h$. Note that the lowest bid in this bin is at least half of any other bid in the bin. If the optimal single-price auction chose, as its selling price, the price of the lowest bid in this bin, then the contribution of each bid in this bin to the revenue is at least half of the bid's value. Since the bin sums to more than $\mathcal{T}/\log h$, this means that the revenue is greater than $\mathcal{T}/(2\log h)$. Thus the optimal fixed pricing can always achieve a revenue of at least $\mathcal{T}/(2\log h)$. ■

One can make this bound strongly polynomial as suggested by Satish Rao and Eva Tardos.

COROLLARY 4.1. $\mathcal{F} \geq \mathcal{T}/(4\log n)$.

Proof. Let v be the optimal price; clearly $v \geq h/n$. If one drops all bids with values below v, $\mathcal{F}$ does not change and $\mathcal{T}$ decreases by at most a factor of two. After the bids are dropped, the ratio of the highest and the lowest bid values is at most n, yielding the desired result. ■

Now we turn our attention to truthful auctions.

5 Generalized Truthful Auction Mechanisms.

By making observations about auctions that encourage utility value bids, in particular the Vickrey auction, we can design general auction mechanisms. We present two general auction mechanisms that facilitate the design of truthful auctions. The first, the *bid-independent auction mechanism*, is based on the observation that the price that a bid is satisfied at must be independent from that bid's actual value. The second, the *random sampling auction mechanism*, is based on the observation that rejected bids can be used to set prices for bids that are to be satisfied.

5.1 Bid-Independent Auctions. The first general truthful auction mechanism that we discuss is one that is typically multi-price, although some variants are single-price. The motivation for this mechanism is the observation that bidder i's bid value should only determine whether bidder i wins or loses the auction (as a threshold). The bid value should not determine bidder i's price.

Let B be the set of all bids and let B_i be the set of bids without bidder i's bid. A bid-independent auction uses a predetermined function, f, from sets of bid values to prices. The auction works as follows. Bidder i wins the auction at price $f(B_i)$ if $b_i \geq f(B_i)$. Otherwise, the bidder loses.

We can immediately see how this generalizes the traditional Vickrey auction. The 1-item Vickrey auction fits into this general framework with $f = \max$. Note that bidder n wins this auction and pays b_{n-1}. If f is the function that returns the kth highest element of the set of bids, we get the k-item Vickrey auction.

5.2 Random Sampling Auctions. Another general truthful auction mechanism is based on randomized sampling. We select a subset B' of B at random, independent of the bid values. Let m be the size of B'. We then compute a function on these sampled bids, $f(B')$, and use this as a threshold value for the $n - m$ bids in the non-sample, $B \setminus B'$. Note that this auction mechanism is inherently single-price.

If a multi-price auction is acceptable, then this auction can be modified to be dual-price with $m \approx n/2$, by using $f(B')$ to compute the threshold to use for bids in $B \setminus B'$ and $f(B \setminus B')$ as the threshold to use for bids in B'. This is a good way to avoid revenue loss due to the rejected bids in the sample; however, it is at the expense of making the auction dual-price.

6 Random Sampling Optimal Threshold Auction.

The random sampling optimal threshold auction takes $f = \mathrm{opt}$, the optimal threshold function, in the random sampling auction. Intuitively we use the sample to get an idea for a good threshold value, then we apply that threshold to the remaining bids. The auction samples $m = n/2$ bids at random, computes opt of this sample, and uses this value as a threshold for the non-sample, accepting all bids above this threshold at the threshold value. In this section we show that the expected revenue of this auction is within a constant factor of $\mathcal{F}$, assuming $\mathcal{F}/h$ is not too small.

For the purpose of simplifying our analysis, we will be analyzing a different method of sampling, one that selects a bid to be in the sample independently at random with probability $1/2$. This method of sampling is simpler to analyze, and it does worse than the former (this can easily be seen when the probabilistic bounds are discussed).

In practice, for the single-price auction, we might want to set $m = n/10$ or even $m = \sqrt{n}$. For the dual-price version of the auction, $m = n/2$ is a good choice.

6.1 Performance Analysis. In this section we show that, under certain assumptions, the expected revenue of the random sampling optimal threshold auction is within a constant factor of $\mathcal{F}$. This result implies that restricting a single-price auction to be truthful does not affect performance by more that a constant factor.

Our analysis of the random sampling auction uses the following lemma, which is a variation of the Chernoff bound (see e.g. [2, 11]).

LEMMA 6.1. *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.

Let t be the number of points already selected when a_{i+1} is considered. 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. ∎

Let $\mathcal{R}$ be the revenue of the random sampling optimal threshold auction. The following theorem shows that $\mathcal{R} = \Omega(\mathcal{F})$ with probability going to one as α goes to ∞.

THEOREM 6.1. *Assume $\alpha h \leq \mathcal{F}$. Then $\mathcal{R} \geq \mathcal{F}/6$ with probability of at least $1 - e^{-\alpha/36} - 40e^{-\alpha/72}$.*

Proof. Let k be the number of bids satisfied in the optimal single-price solution. Consider the optimal fixed pricing revenue of the sample, $\mathcal{F}'$. Since we expect $k/2$ of these bids to be in the sample, $\mathbf{E}[\mathcal{F}'] \geq \mathcal{F}/2$. Applying Lemma 6.1 with A the set of all bids, B the set of bids in the optimal threshold solution on A, and $\delta = 1/3$, we conclude that $|S \cap B| < |B|/3 = k/3$ with probability at least $1 - e^{-k/36}$. The assumption $\alpha h \leq \mathcal{F}$ implies $k \geq \alpha$. Thus $\mathcal{F}' \geq \mathcal{F}/3$ with probability at least $1 - e^{-\alpha/36}$.

Let k_s be the number of bids satisfied by optimal fixed pricing on the sample (i.e. the number of bids in the sample that are at least $\text{opt}(B')$) and let k_n be the number of bids in the non-sample that are at least $\text{opt}(B')$. If $\mathcal{F}' \geq \mathcal{F}/3$ then $k_s \geq \alpha/3$.

Now, assuming that $k_s \geq \alpha/3$, we show that the probability that $k_n < k_s/2$ is small. Note that $k_n < k_s/2$ implies that among the top $(3/2)k_s$ bids, at least k_s are in the sample. Note that the sample and the non-sample are symmetric: taking a random subset containing half of the elements in A is equivalent to taking a complement of a random subset of half of the elements. We apply Lemma 6.1 with S being the non-sample, B being the top $i = (3/2)k_s$ bids, and $\delta = 1/3$, and conclude that the probability that $k_n < k_s/2 = i/3$ is at most $e^{-i/36}$.

If $\mathcal{R} < \mathcal{F}'/2$ then $k_n < k_s/2$ and thus for some $i \geq (3/2)\alpha/3 = \alpha/2$ we have $k_n < i/3$. Using the union bound, the probability that this happens is at most

$$\sum_{i=\alpha/2}^{\infty} e^{-i/36} < 40e^{-\alpha/72}.$$

Using the union bound for the probabilities that $\mathcal{R} \geq \mathcal{F}'/2$ and $\mathcal{F}' \geq \mathcal{F}/3$, we conclude that $\mathcal{R} \geq \mathcal{F}/6$ with probability at least $1 - e^{-\alpha/36} - 40e^{-\alpha/72}$. ∎

Note that in the above theorem, we can trade off the bound on $\mathcal{R}$ and the probability that this bound holds. In particular, for any constant $\epsilon > 0$, we can show that $\mathcal{R} \geq \mathcal{F}/(2+\epsilon)$ with probability that goes to 1 as α goes to infinity, but the convergence is slower for smaller ϵ.

By symmetry, the expected revenue of the dual-price variant of the random sampling auction is twice the expected revenue of the original. One can show that the expected revenue of the dual-price auction is at least $\mathcal{F}/(1+\epsilon)$ if α is large enough.

7 Weighted Pairing.

All truthful auctions we have introduced so far are either single-price or dual-price. In this section we describe a multi-price auction. The *weighted pairing* auction we present is in the bid-independent class with the function f defined as follows:

$$f(B) = b \in B \text{ with probability } \frac{b}{\sum_{b' \in B} b'}$$

Thus, to determine if bidder i wins the auction and at what price, pick a bid $b \in B_i$ with probability proportional to the value of b, i.e., $b/(\mathcal{T} - b_i)$. This pairs b_i with b. If $b \leq b_i$, bidder i wins at cost b, otherwise i loses.

We omit discussion and analysis of the weighted pairing auction and state the two main results. For the (non-trivial) proofs of these results, see the full paper.

THEOREM 7.1. *If $4h \leq \mathcal{T}$, then for the weighted pairing auction $\mathbf{E}[\mathcal{R}] = \Omega(\mathcal{T}/\log h)$.*

THEOREM 7.2. *If $2h \leq F$, then $\mathbf{E}[\mathcal{R}] = \Omega(\mathcal{F}/\sqrt{\log h})$ and this bound is tight.*

Although the weighted pairing auction is not competitive in the worst case, it is only a factor of $\sqrt{\log h}$ away from being competitive. This auction is quite different from our random sampling auctions, and performs relatively well on inputs that are bad for the latter. One may be able to improve this auction; see Section 13.

8 An Upper Bound.

A natural question to ask is if we can improve the bound of Theorem 7.1. We already know that no single-price auction can do better than $\mathcal{F}$. In this section we prove that no truthful multi-price auction can have an expected revenue greater than $\mathcal{F}$. This result augments that of Section 6, where we showed that the performance of a truthful single-price auction is within a constant factor of the optimal single-price auction. Results of this section imply that no truthful multi-price auction performs within a constant factor of the optimal multi-price auction.

Consider a collection of bids $B = \{b_1, \ldots, b_n\}$ with $b_i \leq b_{i+1}$. Note that $b_1 = 1$ and $b_n = h$. Define the following quantities, which are dependent on the auction mechanism:

p_i	the probability a bid i is satisfied,
c_i	expected cost to winning bidder i,
g_i	expected profit (gain) for bidder i.

With probability p_i a bidder i wins. The bidder's expected gain, having won, is their utility value minus the expected price they paid, i.e.

$$g_i = p_i(u_i - c_i) \tag{8.1}$$

The following lemma shows that in a truthful auction, probabilities of winning are monotone functions of bid values.

LEMMA 8.1. *Suppose in a truthful auction $b_i < b_j$. Then $p_i \leq p_j$.*

We omit the proof. The main result of this section is as follows.

THEOREM 8.1. *For any truthful auction, $\mathbf{E}[\mathcal{R}] \leq \mathcal{F}$.*

Proof. In a truthful auction, if a bidder $i-1$ had the utility value of b_i but bids b_{i-1}, their gain would not exceed g_i, thus

$$g_i \geq p_{i-1}(b_i - c_{i-1}). \tag{8.2}$$

So,

$$\begin{aligned} g_i &\geq p_{i-1}(b_i - b_{i-1} + b_{i-1} - c_{i-1}) \\ &= p_{i-1}(b_i - b_{i-1}) + p_{i-1}(b_{i-1} - c_{i-1}) \\ &= p_{i-1}(b_i - b_{i-1}) + g_{i-1}. \end{aligned}$$

We can recursively expand g_{i-1} in the same way until we get to g_1 which is 0 because all bids are satisfied at value at least 1, and get

$$g_i \geq \sum_{j=1}^{i-1} p_j(b_{j+1} - b_j). \tag{8.3}$$

Now let $\mathcal{R}_i$ be the total expected revenue from bidder i. That is

$$\mathcal{R}_i = p_i c_i.$$

We can rearrange equation (8.1) as $p_i c_i = p_i b_i - g_i$ and get

$$\mathcal{R}_i = p_i b_i - g_i.$$

Using equation (8.3) we get

$$\mathcal{R}_i \le p_i b_i - \sum_{j=1}^{i-1} p_j(b_{j+1} - b_j).$$

Looking at the sum of the $\mathcal{R}_i$'s, we see that the first term is mostly canceled by the summation term and we get a telescoping effect.

$$\mathbf{E}[\mathcal{R}] = \sum_{i=1}^{n} \mathcal{R}_i \le \sum_{i=1}^{n} p_i b_i - \sum_{i=1}^{n}\left[\sum_{j=1}^{i-1} p_j(b_{j+1} - b_j)\right].$$

By counting the number of times each $p_j(b_{j+1} - b_j)$ occurs, we can rearrange the second summation to get

$$\mathbf{E}[\mathcal{R}] \le \sum_{i=1}^{n} p_i b_i - \sum_{j=1}^{n-1} p_j(b_{j+1} - b_j)(n-j)$$

$$\mathbf{E}[\mathcal{R}] \le p_n b_n + \sum_{j=1}^{n-1} p_i b_i - \sum_{j=1}^{n-1} p_j(b_{j+1} - b_j)(n-j).$$

$$= p_n b_n + \sum_{j=1}^{n-1} p_j\left[b_j - (b_{j+1} - b_j)(n-j)\right].$$

Regrouping

$$\mathbf{E}[\mathcal{R}] \le p_n b_n + \sum_{j=1}^{n-1} p_j\left[b_j(n-j+1) - b_{j+1}(n-j)\right].$$

Now, let $V_j = b_j(n-j+1)$. Intuitively, this is the revenue attained by using b_j as the sale price in fixed pricing. Note that $V_j \le \mathcal{F}$.

$$\mathbf{E}[\mathcal{R}] \le p_n V_n + \sum_{j=1}^{n-1} p_j\left(V_j - V_{j+1}\right).$$

Rearrange this sum to sum over V_j instead of p_j and for symmetry, define $p_0 = 0$.

$$\mathbf{E}[\mathcal{R}] \le \sum_{j=1}^{n} \left(p_j - p_{j-1}\right) V_j.$$

But, $V_j \le \mathcal{F}$ and by Lemma 8.1, $p_j - p_{j-1}$ is non-negative.

$$\mathbf{E}[\mathcal{R}] \le \mathcal{F}\sum_{j=1}^{n}(p_j - p_{j-1})$$

This sum telescopes to $p_n - p_0$ and $p_0 = 0$ so we have

$$\mathbf{E}[\mathcal{R}] \le p_n \mathcal{F} \le \mathcal{F}.$$

■

9 Deterministic Auctions.

We have shown several randomized auctions with an $\Omega(\mathcal{F})$ expected performance under the assumption that $\alpha h \le \mathcal{F}$ and α is big enough. In this section we study deterministic auctions.

The *deterministic optimal threshold auction* is the bid-independent auction with $f = \text{opt}$, the optimal threshold function defined in Section 4. The only difference between the deterministic optimal threshold auction and the optimal fixed pricing mechanism is that the former uses threshold $\text{opt}(B_i)$ for bidder i and the latter uses $\text{opt}(B)$. Recall that B_i and B only differ in that b_i is not in B_i. From this, we might also expect that for large n with suitable constraints on h, the deterministic optimal threshold auction would perform to within a constant fraction of $\mathcal{F}$. As we will see shortly, this is not the case. An interesting result that we will not show in this paper is that the deterministic optimal threshold auction is single-price.

9.1 Upper Bound for Deterministic Bid-Independent Auctions.

In this section we prove the following upper bound:

THEOREM 9.1. *For any truthful deterministic bid-independent auction and any constant α, there exists an input for which $\mathcal{R}/\mathcal{F} = O(1/h)$ and $\alpha h \le \mathcal{F}$.*

Proof. Let f be the (deterministic) function that defines the auction. Consider a bipolar input with n_h bids at value h and n_ℓ bids at value 1. Restricted to such inputs, f is a function of n_h and n_ℓ. Note that we can assume, without loss of generality, that f takes on only two values, 1 and h. Other values of f lead to smaller revenues.

We wish to find an input family such that the bid-independent auction with function f has revenue $O(\mathcal{F}/h)$. To obtain this family, we chose n_h and n_ℓ in such a way that $f(n_h - 1, n_\ell) = 1$, $f(n_h, n_\ell - 1) = h$, and $n_h > \alpha$ (implying $\alpha h \le \mathcal{F}$). For such an input, $\mathcal{R} = n_h$ and $\mathcal{F} \ge h n_h$ so $\mathcal{R}/\mathcal{F} \le 1/h$. Our goal now is, given a deterministic f, to find values of n_h and n_ℓ that have the above properties.

Consider the n_h, n_ℓ plane. For a fixed m look at the line $n_h = k$ and $n_\ell = m - k$, and consider the line segment connecting $(0, m)$ and $(m, 0)$. We need to find a value of k with $k > \alpha$ where, when k increases by one, f changes from 1 to h.

Set $m = h^2\alpha$. Assume that $f(\alpha, m-\alpha) = 1$. As we increase k from $k = \alpha$ the value of $f(k, m-k)$ must change from 1 to h because for $k = m$ we have $f(k, m-k) = f(m, 0) = h$. Thus it must be at some k^* that $f(k^*, m-k^*) = h$ and $f(k^*-1, m-k^*+1) = 1$. If we now choose $n_h = k^*$ and $n_\ell = m - k^* + 1$, we satisfy our criteria that $n_h > \alpha$ and that $f(n_h - 1, n_\ell) = 1$ and $f(n_h, n_\ell - 1) = h$. Thus, $\mathcal{R}/\mathcal{F} = O(1/h)$.

Suppose now that our assumption that $f(\alpha, m - \alpha) = 1$ is false and instead it is h. Then for $n_h = k = \alpha$ and $n_\ell = m - k + 1$ we have $\mathcal{F} = m + 1 = h^2\alpha + 1$ and $\mathcal{R} \leq h\alpha$ so

$$\mathcal{R}/\mathcal{F} \leq \frac{h\alpha}{h^2\alpha+1} \leq \frac{h\alpha}{h^2\alpha} = \frac{1}{h}.$$

Thus for any bid-independent auction with deterministic function f there exists a distribution such that $\mathcal{R} = O(\mathcal{F}/h)$. ■

9.2 Truthful Deterministic Auctions are Bid-Independent. The following lemma allows us to extend the result of Theorem 9.1 to arbitrary deterministic auctions.

LEMMA 9.1. *Any truthful deterministic auction is bid-independent.*

See the full version of the paper for the proof. Using Lemma 9.1, we generalize Theorem 9.1 as follows.

THEOREM 9.2. *For any truthful deterministic auction and any constant α, there exists an input for which $\mathcal{R}/\mathcal{F} = O(1/h)$ and $\alpha h \leq \mathcal{F}$.*

In terms of asymptotic worst-case performance, deterministic auctions are significantly worse than randomized auctions. This is not to say that deterministic auctions are bad to use for all input families. In fact our experimental results reveal that for many families, the deterministic optimal threshold auction works very well. Adequate knowledge of the bidding distribution may make it possible to use a deterministic auction.

10 Experimental Results.

Our theoretical analysis of auctions has limitations. Worst-case analysis against $\mathcal{F}$ or $\mathcal{T}/\log h$ leaves a $\log h$ gap for inputs which come from "typical" (as opposed to tailored to be hard) distributions. In addition, constant factors we obtain in our analysis are often too pessimistic. Theoretical analysis for specific distributions seems quite difficult even for simple distributions.

In practice, constant factors of the auction revenue are important. We introduced several auction mechanisms that provably perform within a constant factor of each other in worst-case. However, we do not know which one is better. We would also like to know how these auctions compare to fixed pricing with imperfect market analysis.

We turn to experiments to answer these questions. In our experiments, we simulate various auctions on several input families and see how they compare. Below we present experimental results for two problem families. See the complete paper for more experimental results.

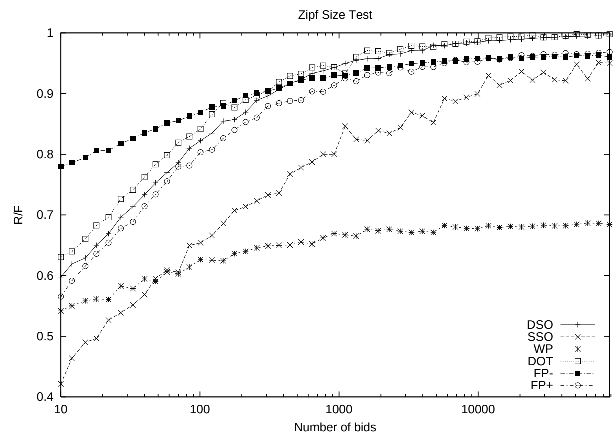

Figure 1: Zipf(theta=0.5, high=n) with $n = [10, 100k]$

In the first problem family, each bid is chosen independently from the Zipf distribution [18] with $\theta = 1/2$. This is a generalization of the distribution with 80% of the total bid value coming form 20% of the bidders. For i in the interval $[1, high]$, we define $\mathbf{Pr}[X = i] = c/i^\theta$, with c chosen so that the probabilities integrate to one. Figure 1 shows the results of simulation for the Zipf family as the number of bidders varies between 10 and 100,000. This family has the property that any uniformly chosen random subset of the bids has the same distribution as the original. Because of this property, the random sampling auctions perform very well for these families. In the plot ratios $\mathcal{R}/\mathcal{F}$ are reported for the following mechanisms:

DSO	dual-price sampling optimal threshold.
SSO	single-price sampling optimal threshold, $m = \sqrt{n}$.
WP	weighted pairing.
DOT	deterministic optimal threshold.
FP−	fixed pricing with optimal price less 25%.
FP+	fixed pricing with optimal price plus 25%.

For large n, DSO and DOT are the best auctions. As n increases, the ratio of their revenue to $\mathcal{F}$ approaches 1, whereas the ratios for FP− and FP+ approach a constant less than 1. As a result, even for moderately large n our best auctions perform better than fixed pricing with a 25% price error. On average-case distributions, WP is the worst auction; its ratio asymptotically approaches 2/3.

In the bipolar family, bids have only two values, 1 and h, and the ratio of high bids to the total number of bids varies. The results appear in Figure 2.

There are several key things to note about the bipolar family. First, it demonstrates the problem with

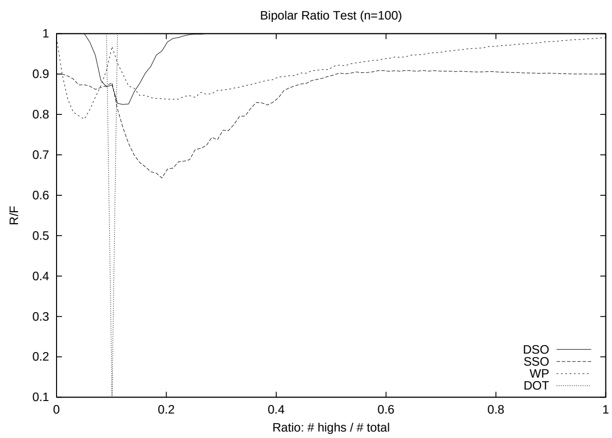

Figure 2: Bipolar(low=1,high=10,ratio=$[0, 1]$)

the DOT auction. There is a sharp dip in the revenue of DOT precisely when the number of high bids is 10 and the number of low bids is 90. This is because $\text{opt}(B \setminus \{h\}) = 1$ and $\text{opt}(B \setminus \{1\}) = h$ in this scenario and thus 10 high bids get satisfied at price 1 and the 90 low bids get rejected. The optimal solution in this case is to accept the 10 high bids at price $h = 10$. Aside from this case, DOT performs well. The randomized optimal threshold variants have degraded revenues for the number of high bids around 10. However, due to randomness, the revenue loss is significantly smaller, but spread over a wider region. The DSO auction outperforms WP and SSO for most ratios. The WP auction usually outperforms SSO.

11 Bounded Supply.

Up to this point, we have studied the unlimited supply case which is motivated by the digital goods market where the cost of making a copy of an item is negligible. In this section we consider the case where the number of items available for sale is bounded. This case is typical of physical goods markets. We denote the number of items available by k. Here again, the seller wishes to maximize their revenue, possibly not selling all of the items. Note that the definitions of truthful and competitive auctions, which we stated for the unlimited supply case, also apply to the bounded supply case. We denote by $\mathcal{F}_k$ the revenue for optimal fixed pricing that sells at most k items, and it is this quantity that we wish to be competitive with.

The bounded supply case is a generalization of the unlimited supply case as items are available in unlimited supply when the number of available items is the same as the number of bidders (i.e. $k = n$). Where unlimited supply is one extreme of the bounded supply case, *scarce* supply is another extreme. In the scarce supply case, the optimal fixed pricing revenue is maximized by selling all the items when the number of available items is around k. Previous work on auctions concentrated on the scarce supply case. The multi-item Vickrey auction is the best single-price auction for the scarce supply case (and it is competitive). Our results, extended to bounded supply, are competitive in the full range of the supply-demand spectrum with only the assumption that $\alpha h < \mathcal{F}_k$.

We now show how to extend our optimal threshold sampling auctions to the k-item bounded supply case. Let opt_k be the function that, given a set of bids, returns the optimal threshold that sells k items or less. This function, on the entire set of bids, gives the threshold to use for $\mathcal{F}_k$. The single-price sampling optimal threshold mechanisms can now be modified to use threshold function $\text{opt}_{mk/(n-m)}$ on sample of size m. If this results in too many bids being satisfied, arbitrarily (e.g. at random) reject bids until there are only k left. One can show that with high probability, the number of bids rejected will be small and that this auction is competitive.

In the dual-price auction with sample size $m = n/2$, use $\text{opt}_{k/2}$ so that about $k/2$ bids are selected from each of the sample and the non-sample. Once can show that the resulting auction is also truthful and competitive.

We have shown that in a relatively straight-forward way, our sampling auctions extend from the unlimited to the bounded supply case. One can extend the deterministic optimal threshold auction to bounded supply as well. Also, all of our upper bounds apply to bounded supply because it generalizes unlimited supply.

To generalize Theorem 4.1 for the k-item case, recall the alternative definition of $\mathcal{T}$ as the revenue due to the optimal untruthful multi-price auction. Define $\mathcal{T}_k$ to be the revenue of the optimal multi-price auction restricted to only satisfying k bidders (i.e. the sum of the highest k bids) and $\mathcal{F}_k$ as above. Then the generalized result is $\mathcal{F}_k \geq \mathcal{T}_k/(2 \log h)$.

One result we do not know how to extend to bounded supply is that for the weighted pairing auction.

12 Concluding Remarks.

We have demonstrated that there exist truthful auctions for unlimited supply markets. We have shown randomized auctions that are competitive in that they yield revenue that is within a constant factor of optimal fixed pricing. We have shown that this result is tight up to a constant factor, even for multi-price auctions. We have also shown that no deterministic auction is competitive in the worst-case. Finally, via simulations, we have argued that our auctions compare favorably to fixed pricing with market analysis.

For unlimited supply markets, our analysis assumes that there is no cost for producing the items being auctioned. With the following modification we can also accommodate non-zero marginal costs. If the marginal cost is v per item, then first subtract v from each bid and reject all negative value bids. After running the auction, add v back to the selling price of all winning bids. If the marginal cost of producing k items is a more complicated function of k, we can modify the opt function used in the sampling optimal threshold auction to take into account these marginal costs. In this case the opt function would, as above for bounded supply, need to be parameterized by the ratio of the sample size and the non-sample size so as to correctly use marginal cost information.

To prevent cheating by the auctioneer or the bidders, one may need a trusted third party or a special cryptographic protocol. Note that cheating prevention is a problem shared by all on-line auctions. Related results appear in [5].

13 Open Problems.

We proved that the weighted pairing auction is not competitive. It is possible, however, that a variant of this auction that uses different weighting is competitive.

The weighted pairing auction can have revenue greater than $\mathcal{F}$ for some random pairings. It seems like this might not be the case for the dual-price sampling optimal threshold auction. Can one prove that it is always the case that for any randomly chosen sample, the auction revenue does not exceed $\mathcal{F}$?

We have not considered issues such as the extent to which bidders can remain anonymous or bid values can remain secret. See [7] for ways to maintain bid secrecy in an on-line Vickrey auction.

A significant issue in auctions like ours is resistance to adversarial attacks. How resistant are our auctions to bidder collusion, and can collusion resistance be improved? How well do our auctions resist attacks such as a competitor attempting to reduce the revenue of an auction by submitting a large number of low bids?

Repeated auctions for the same item may be of interest in some applications. In this case, the challenge is to design a truthful auction mechanism.

Acknowledgements

We would like to thank the following people for their contributions and useful discussions relating to this work: Amos Fiat, Stuart Haber, Jim Horning, Anna Karlin, Umesh Maheshwari, Satish Rao, Eva Tardos, Bob Tarjan, Radek Vingralek, and Stephen Weeks.

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. 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.

[5] M. K. Franklin and M. K. Reiter. The design and implementation of a secure auction service. *IEEE Trans. on Sftware Eng.*, 22:302–312, 1996.

[6] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.

[7] M. Harkaby, J. D. Tygar, and H. Kikuchi. Electronic Auctions with Private Bids. In *Proc. of the 3rd USENIX Workshop on Electr. Commerce*, pages 61–73. USENIX Assoc., Berkeley, CA, 1998.

[8] 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.

[9] N. Lineal. Game Theoretic Aspects of Computing. In *Handbook of Game Theory*, volume 2, pages 1339–1395. Elseveir Science Publishers B.V., 1994.

[10] A. Mas-Colell, M.D. Whinston, and J. R. Green. *Microeconomic Theory.* Oxford University Press, New York, Oxford, 1995.

[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 Theory of Computing*, pages 129–140. ACM Press, New York, 1999.

[13] M. J. Osborne and A. Rubistein. *A Course in Game Theory.* MIT Press, Cambridge, MA, 1994.

[14] A. E. Roth and M. A. Oliveira Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis.* Cambridge University Press, 1990.

[15] D. D. Sleator and R. E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Communications of the ACM*, 28:202–208, 1985.

[16] M. R. Varian. Economic Mechanism Design for Computerized Agents. In *Proc. of the 1st USENIX Workshop on Electr. Commerce.* USENIX Assoc., Berkeley, CA, 1995.

[17] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.

[18] G. K. Zipf. *Human Behavior and the Principle of Least Effort, an Introduciton to Human Ecology.* Addison-Wesley, 1949.