

A Lower Bound on the Competitive Ratio of Truthful Auctions

Andrew V. Goldberg¹, Jason D. Hartline¹, Anna R. Karlin², and Michael Saks³

¹ Microsoft Research, SVC/5, 1065 La Avenida, Mountain View, CA 94043.*
{goldberg, hartline}@microsoft.com

² Computer Science Department, University of Washington.
karlin@cs.washington.edu

³ Department of Mathematics–Hill Center, Rutgers University. 110 Frelinghuysen Rd., Piscataway, NJ 08854. ** saks@math.rutgers.edu

Abstract We study a class of single-round, sealed-bid auctions for a set of identical items. We adopt the worst case competitive framework defined by [1,2] that compares the profit of an auction to that of an optimal single price sale to at least two bidders. In this framework, we give a lower bound of 2.42 (an improvement from the bound of 2 given in [2]) on the competitive ratio of any truthful auction, one where each bidder's best strategy is to declare the true maximum value an item is worth to them. This result contrasts with the 3.39 competitive ratio of the best known truthful auction [3].

1 Introduction

A combination of recent economic and computational trends, such as the negligible cost of duplicating digital goods and, most importantly, the emergence of the Internet as one of the most important arenas for resource sharing between parties with diverse and selfish interests, has created a number of new and interesting dynamic pricing problems. It has also cast new light on more traditional problems such as the problem of profit maximization for the seller in an auction.

A number of recent papers [1,2,3] have considered the problem of designing auctions, for selling identical units of an item, that perform well in *worst case* under unknown market conditions. In these auctions, there is a seller with ℓ units for sale, and bidders each interested in obtaining one of them. Each bidder has a valuation representing how much the item is worth to them. The auction is performed by soliciting a sealed bid from each of the bidders, and deciding on the allocation of units to bidders and the prices to be paid by the bidders. The bidders are assumed to follow the strategy of bidding so as to maximize their personal utility, the difference between their valuation and the price they pay. To handle the problem of designing and analyzing auctions where bidders may

* Work was done while second author was at the University of Washington.

** Supported by NSF grant CCR-9988526. Part of the work was done while visiting Microsoft Research.

falsely declare their valuations to get a better deal, we will adopt the solution concept of *truthful mechanism design* (see, e.g., [1,4,5]). In a truthful auction, truth-telling, i.e., revealing their true valuation as their bid, is an optimal strategy for each bidder regardless of the bids of the other bidders. In this paper, we will restrict our attention to truthful (a.k.a., incentive compatible or strategyproof) auctions.

In research on such auctions, a form of *competitive analysis* is used to gauge auction revenue. Specifically, a truthful auction’s performance on a particular bid vector is evaluated by comparing it against the profit that could be achieved by an “optimal” omniscient auction, one that knows the true valuations of the bidders in advance. An auction is β -*competitive* if it achieves a profit that is within a factor of $\beta \geq 1$ of optimal *on every input*. The goal then becomes to design the auction with the best competitive ratio, i.e., the auction that is β -competitive with the smallest possible value of β .

A particularly interesting special case of the auction problem is the *unlimited supply* case. In this case the number of units for sale is at least the number of bidders in the auction. This is natural for the sale of digital goods where there is negligible cost for duplicating and distributing the good. Pay-per-view television and downloadable audio files are examples of such goods.

For the unlimited supply auction problem, the competitive framework introduced in [1] and further refined in [2] uses the profit of the *optimal omniscient single priced mechanism that sells at least two units* as the benchmark for competitive analysis. The assumption that two or more units are sold is necessary because in the worst case it is impossible to obtain a constant fraction of the profit of the optimal mechanism when it sells only one unit [1]. In this worst case competitive framework, the best known auction for the unlimited supply has a competitive ratio of 3.39 [3].

In this paper we also consider the case where the number of units for sale, ℓ , is limited, i.e., less than the number of bidders. At the opposite extreme from unlimited supply, is the limited supply case with $\ell = 2$.⁴ In this case the Vickrey auction [4], which sells to the highest bidder at the second highest bid value, obtains the optimal worst case competitive ratio of 2 [2].

The main result of this paper is a lower bound on the competitive ratio of any randomized auction. For $\ell = 2$, this lower bound is 2 (this was originally proven in [2], though we give a much simpler proof of it here). For $\ell = 3$, the lower bound is $13/6 \approx 2.17$, and as ℓ grows the bound approaches 2.42 in the limit. We conjecture that this lower bound is tight. Yet, even in the case of three units, the problem of constructing the auction matching our lower bound of $13/6$ is open.

The rest of the paper is organized as follows. In Section 2 we give the mathematical formulation of the auction problem that we will be studying, and we describe the competitive framework that is used to analyze such auctions in worst case. In Section 3 we give our main result, a bound on how well any auc-

⁴ Notice that the competitive framework is not well defined for the $\ell = 1$ case as the optimal auction that sells at least two units cannot sell just one unit.

tion can perform in worst case. In Section 4 we describe attempts to obtain a matching upper bound.

2 Preliminaries and Notation

We consider single-round, sealed-bid auctions for a set of ℓ identical units. As mentioned in the introduction, we adopt the game theoretic solution concept of truthful mechanism design. A useful simplification of the problem of designing truthful auctions is obtained through the following algorithmic characterization. Related formulations to the one we give here have appeared in numerous places in recent literature (e.g., [6,7,2,8]). To the best of our knowledge, the earliest dates back to the 1970s [9].

Definition 1. *Given a bid vector of n bids, $\mathbf{b} = (b_1, \dots, b_n)$, let \mathbf{b}_{-i} denote the vector of with b_i replaced with a ‘?’ , i.e.,*

$$\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, ?, b_{i+1}, \dots, b_n).$$

Definition 2 (Bid-independent Auction, BI_f). *Let f be a function from bid vectors (with a ‘?’) to prices (non-negative real numbers). The deterministic bid-independent auction defined by f , BI_f , works as follows. For each bidder i :*

1. Set $t_i = f(\mathbf{b}_{-i})$.
2. If $t_i < b_i$, bidder i wins at price t_i
3. If $t_i > b_i$, bidder i loses.
4. Otherwise, ($t_i = b_i$) the auction can either accept the bid at price t_i or reject it.

A randomized bid-independent auction is a distribution over deterministic bid-independent auctions.

The proof of the following theorem can be found, for example, in [2].

Theorem 1. *An auction is truthful if and only if it is equivalent to a bid-independent auction.*

Given this equivalence, we will use the the terminology *bid-independent* and *truthful* interchangeably. We denote the profit of a truthful auction \mathcal{A} on input \mathbf{b} as $\mathcal{A}(\mathbf{b})$. This profit is given by the sum of the prices charged bidders that are not rejected. For a randomized bid-independent auction, $\mathcal{A}(\mathbf{b})$ and $f(\mathbf{b}_{-i})$ are random variables.

It is natural to consider a worst case competitive analysis of truthful auctions. In the competitive framework of [2] and subsequent papers, the performance of a truthful auction is gauged in comparison to the *optimal auction that sells at least two units*. There are a number reasons to choose this metric for comparison, interested readers should see [2] or [10] for a more detailed discussion.

Definition 3. The optimal single price omniscient auction that sells at least two units (and at most ℓ units), $\mathcal{F}^{(2,\ell)}$, is defined as follows: Let \mathbf{b} be a bid vector of n bids, and let v_i be the i -th largest bid in the vector \mathbf{b} . Auction $\mathcal{F}^{(2,\ell)}$ on \mathbf{b} chooses $k \in \{2, \dots, \ell\}$ to maximize kv_k . The k highest bidders are each sold a unit at price v_k (ties broken arbitrarily); all remaining bidders lose. Its profit is:

$$\mathcal{F}^{(2,\ell)}(\mathbf{b}) = \max_{2 \leq k \leq \ell} kv_k.$$

In the unlimited supply case, i.e., when $\ell = n$, we define $\mathcal{F}^{(2)} = \mathcal{F}^{(2,n)}$.

Definition 4. We say that auction \mathcal{A} is β -competitive if for all bid vectors \mathbf{b} , the expected profit of \mathcal{A} on \mathbf{b} satisfies

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \geq \frac{\mathcal{F}^{(2,\ell)}(\mathbf{b})}{\beta}.$$

The competitive ratio of the auction \mathcal{A} is the infimum of β for which the auction is β -competitive.

2.1 Limited Supply Versus Unlimited Supply

Throughout the remainder of this paper we will be making the assumption that $n = \ell$, i.e., the number of bidders is equal to the number of items for sale. The justification for this is that any lower bound that applies to the $n = \ell$ case also extends to the case where $n \geq \ell$. To see this, note that an ℓ item auction \mathcal{A} that is β -competitive for any $n > \ell$ bidder input must also be β -competitive on the subset of all n bidder bid vectors that have $n - \ell$ bids at value zero. Thus, we can simply construct an \mathcal{A}' that takes ℓ bidder input \mathbf{b}' , augments it with $n - \ell$ zeros to get \mathbf{b} , and simulates the outcome of \mathcal{A} on \mathbf{b} . Since $\mathcal{F}^{(2)}(\mathbf{b}') = \mathcal{F}^{(2,\ell)}(\mathbf{b})$, \mathcal{A}' obtains at least the competitive ratio of \mathcal{A} .

In the other direction, a reduction from the unlimited supply auction problem to the limited supply auction problem given in [10] shows how to take an unlimited supply auction that is β -competitive with $\mathcal{F}^{(2)}$ and construct a limited supply auction parameterized by ℓ that is β -competitive with $\mathcal{F}^{(2,\ell)}$.

Henceforth, we will assume that we are in the unlimited supply case, and we will examine lower bounds for limited supply problems by placing a restriction on the number of bidders in the auction.

2.2 Symmetric Auctions

In the remainder of this paper, we restrict attention to *symmetric auctions*. An auction is symmetric if its output is not a function of the order of the bids in the input vector, \mathbf{b} . We note that there is no loss of generality in this assumption, as the following result shows.

Lemma 1. For any β -competitive asymmetric truthful auction there is a symmetric randomized truthful auction with competitive ratio at least β .

Proof. Given a β -competitive asymmetric truthful auction, \mathcal{A} , we construct a symmetric truthful auction \mathcal{A}' that first permutes the input bids \mathbf{b} at random to get $\pi(\mathbf{b})$ and then runs \mathcal{A} on $\pi(\mathbf{b})$. Note, $\mathcal{F}^{(2)}(\mathbf{b}) = \mathcal{F}^{(2)}(\pi(\mathbf{b}))$ and since \mathcal{A} is β -competitive on $\pi(\mathbf{b})$ for any choice of π , \mathcal{A}' is β -competitive on \mathbf{b} . ■

2.3 Example: The Vickrey Auction

The classical truthful auction is the 1-item Vickrey auction (a.k.a. the second price auction). This auction sells to the highest bidder at the second highest bid value. To see how this fits into the bid-independent framework, note that the auction BI_{\max} (the bid-independent auction with $f = \max$) does exactly this (assuming that the largest bid is unique).

As an example we consider the competitive ratio of the Vickrey auction in the case where there are only two bidders. Given two bids, $\mathbf{b} = \{b_1, b_2\}$, the optimal single price sale of two units just sells both units for the smaller of the two bid values, i.e., the optimal profit is $\mathcal{F}^{(2)}(\mathbf{b}) = 2 \min(b_1, b_2)$. Of course, the 1-item Vickrey auction sells to the highest bidder at the second highest price and thus has a profit of $\min(b_1, b_2)$. Therefore, we have:

Observation 1 *The Vickrey auction on two bidders is 2-competitive.*

It turns out that this is optimal for two bidders. Along with the general lower bound of 2.42, in the next section we give a simplified proof of the result, originally from [2], that no two bidder truthful auction is better than 2-competitive.

3 A Lower Bound on the Competitive Ratio

In this section we prove a lower bound on the competitive ratio of any truthful auction in comparison to $\mathcal{F}^{(2)}$; we show that for any randomized truthful auction, \mathcal{A} , there exists an input bid vector, \mathbf{b} , on which

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \leq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{2.42}.$$

In our lower bound proof we will be considering randomized distributions over bid vectors. To avoid confusion, we will adopt the following notation. A real valued random variable will be given in uppercase, e.g., X and T_i . In accordance with this notation, we will use B_i as the random variable for bidder i 's bid value. A vector of real valued random variables will be a bold uppercase letter, e.g., \mathbf{B} is a vector of random bids.

To prove the lower bound, we analyze the behavior of \mathcal{A} on a bid vector chosen from a probability distribution over bid vectors. The outcome of the auction is then a random variable depending on both the randomness in \mathcal{A} and the randomness in \mathbf{B} . We will give a distribution on bidder bids and show that it satisfies $\mathbf{E}_{\mathbf{B}}[\mathbf{E}_{\mathcal{A}}[\mathcal{A}(\mathbf{B})]] \leq \frac{\mathbf{E}_{\mathbf{B}}[\mathcal{F}^{(2)}(\mathbf{B})]}{2.42}$. We then use the following fact to claim that there must exist a fixed choice of bids, \mathbf{b} (depending on \mathcal{A}), for which $\mathbf{E}[\mathcal{A}(\mathbf{b})] \leq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{2.42}$.

Fact 1 Given random variable X and two functions f and g , $\mathbf{E}[f(X)] \leq \mathbf{E}[g(X)]$ implies that there exists x such that $f(x) \leq g(x)$.

As a quick proof of this fact, observe that if for all x , $f(x) > g(x)$ then it would be the case that $\mathbf{E}[f(X)] > \mathbf{E}[g(X)]$ instead of the other way around.

A key step in obtaining the lower bound is in defining a distribution over bid vectors on which any truthful auction obtains the same expected revenue.

Definition 5. Let the random vector of bids $\mathbf{B}^{(n)}$ be n i.i.d. bids generated from the distribution with each bid B_i satisfying $\Pr[B_i > z] = 1/z$ for all $z \geq 1$.

Lemma 2. For $\mathbf{B}^{(n)}$ defined above, any truthful auction, \mathcal{A} , has expected revenue satisfying,

$$\mathbf{E}[\mathcal{A}(\mathbf{B}^{(n)})] \leq n.$$

Proof. Consider a truthful auction \mathcal{A} . Let T_i be the price offered to bidder i in the bid-independent implementation of \mathcal{A} . T_i is a random variable depending on \mathcal{A} and \mathbf{B}_{-i} and therefore T_i and B_i are independent random variables. Let P_i be the price paid by bidder i , i.e., 0 if $B_i < T_i$ and T_i otherwise. For $t \geq 0$, $\mathbf{E}[P_i \mid T_i = t] = t \cdot \Pr[B_i > t \mid T_i = t] = t \cdot \Pr[B_i > t] \leq 1$, since B_i is independent of T_i . Therefore $\mathbf{E}[P_i] \leq 1$ and $\mathbf{E}[\mathcal{A}(\mathbf{B}^{(n)})] = \sum_i \mathbf{E}[P_i] \leq n$. ■

For the input $\mathbf{B}^{(n)}$ an auction attempting to maximize the profit of the seller has no reason to ever offer prices less than one. The proof of the above lemma shows that any auction that always offers prices of at least one has expected revenue exactly n .

3.1 The $n = 2$ Case

To give an outline for how our main proof will proceed, we first present a proof that the competitive ratio for a two bidder auction is at least 2. Of course, the fact that the 1-item Vickrey auction achieves this competitive ratio means that this result is tight. The proof we give below simplifies the proof of the same result given in [2].

Lemma 3. $\mathbf{E}[\mathcal{F}^{(2)}(\mathbf{B}^{(2)})] = 4$.

Proof. From the definition of $\mathcal{F}^{(2)}$, $\mathcal{F}^{(2)}(\mathbf{B}^{(2)}) = 2 \min \mathbf{B}^{(2)}$. Therefore, for $z \geq 2$, $\Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(2)}) > z] = \Pr[B_1 > z/2 \wedge B_2 > z/2] = 4/z^2$. Using the definition of expectation for non-negative continuous random variables of $\mathbf{E}[X] = \int_0^\infty \Pr[X > x] dx$ we have

$$\mathbf{E}[\mathcal{F}^{(2)}(\mathbf{B}^{(2)})] = 2 + \int_2^\infty (4/z^2) dz = 4. \quad \blacksquare$$

Lemma 4. The optimal competitive ratio for a two bidder auction is 2.

The proof of this lemma follows directly from Lemmas 2 and 3, and Fact 1.

3.2 The General Case

For the general case, as in the two bidder case, we must compute the expectation of $\mathcal{F}^{(2)}(\mathbf{B}^{(n)})$.

Lemma 5. *For n bids from the above distribution, the expected value of $\mathcal{F}^{(2)}$ is*

$$\mathbf{E}\left[\mathcal{F}^{(2)}(\mathbf{B}^{(n)})\right] = n - n \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}.$$

Proof. In this proof we will get a closed form expression for $\Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z]$ and then integrate to obtain the expected value. Note that all bids are at least one and therefore, we will assume that $z \geq n$. Clearly for $z < n$, $\Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z] = 1$. Let V_i be a random variable for the value of the i th largest bid, e.g., $V_1 = \max_i B_i$. To get a formula for $\Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)})]$, we define a recurrence based on the random variable $F_{n,k}$ defined as

$$F_{n,k} = \max_i (k+i)V_i.$$

Intuitively, $F_{n,k}$ represents the optimal single price revenue from $\mathbf{B}^{(n)}$ and an additional k consumers each of which has a value equal to the highest bid, V_1 . To define the recurrence, fix n, k , and z and define the events \mathcal{H}_i for $1 \leq i \leq n$. Intuitively, the event \mathcal{H}_i represents the fact that i bidders in $\mathbf{B}^{(n)}$ and the k additional consumers have bid high enough to equally share z , while no larger set of $j > i$ bidders of $\mathbf{B}^{(n)}$ can do the same.

$$\mathcal{H}_i = V_i \geq z/(k+i) \wedge \bigwedge_{j=i+1}^n V_j < z/(k+j)$$

$$\Pr[\mathcal{H}_i] = \binom{n}{i} \left(\frac{k+i}{z}\right)^i \Pr[F_{n-i,k+i} < z].$$

Note that events \mathcal{H}_i are disjoint and that $F_{n,k}$ is at least z if and only if one of the \mathcal{H}_i occurs. Thus,

$$\begin{aligned} \Pr[F_{n,k} > z] &= \Pr\left[\bigwedge_{i=1}^n \mathcal{H}_i\right] = \sum_{i=1}^n \Pr[\mathcal{H}_i] \\ &= \sum_{i=1}^n \binom{n}{i} \left(\frac{k+i}{z}\right)^i \Pr[F_{n-i,k+i} < z]. \end{aligned} \quad (1)$$

Equation (1) defines a two dimensional recurrence. The base case of this recurrence is given by $F_{0,k} = 0$. We are interested in $\mathcal{F}^{(2)}(\mathbf{B}^{(n)})$ which is the same as $F_{n,0}$ except that we ignore the \mathcal{H}_1 case. This gives

$$\begin{aligned} \Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z] &= \Pr[F_{n,0} > z] - \Pr[\mathcal{H}_1] \\ &= \Pr[F_{n,0} > z] - \frac{n}{z} \Pr[F_{n-1,1} < z]. \end{aligned} \quad (2)$$

To obtain $\Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)})]$ we can solve the recurrence for $F_{n,k}$ given by Equation (1). We will show that the solution is:

$$\Pr[F_{n,k} > z] = 1 - \left(\frac{z-k}{z}\right)^n \left(\frac{z-k-n}{z-k}\right). \quad (3)$$

Note that (3) is correct for $n = 0$. We show that it is true in general inductively. Substituting in our proposed solution (3) into (1) we obtain:

$$\begin{aligned} \Pr[F_{n,k} > z] &= \sum_{i=1}^n \binom{n}{i} \left(\frac{k+i}{z}\right)^i \left(\frac{z-k-i}{z}\right)^{n-i} \left(\frac{z-k-n}{z-k-i}\right) \\ &= \frac{z-k-n}{z^n} \sum_{i=1}^n \binom{n}{i} (k+i)^i (z-k-i)^{n-i-1}. \end{aligned} \quad (4)$$

We now apply the following version of Abel's Identity [11]:

$$\frac{(x+y)^n}{x} = \sum_{j=0}^n \binom{n}{j} (x+j)^{j-1} (y-j)^{n-j}.$$

Making the change of variables, $j = n - i$, $x = z - k - n$, and $y = k + n$ we get:

$$\frac{z^n}{z-k-n} = \sum_{i=0}^n \binom{n}{i} (k+i)^i (z-k-i)^{n-i-1}.$$

We subtract out the $i = 0$ term and plug this identity into (4) to get

$$\begin{aligned} \Pr[F_{n,k} > z] &= \frac{z-k-n}{z^n} \left(\frac{z^n}{z-k-n} - (z-k)^{n-1} \right) \\ &= 1 - \left(\frac{z-k}{z}\right)^n \left(\frac{z-k-n}{z-k}\right). \end{aligned}$$

Thus, our closed form expression for the recurrence is correct.

Recall our goal is to compute $\Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z]$. Equation (3) shows that $\Pr[F_{n,0} > z] = n/z$. This combined with Equation (2) and Equation (3) gives the following for $z \geq n$:

$$\begin{aligned} \Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z] &= \frac{n}{z} - \frac{n}{z} \Pr[F_{n-1,1} < z] \\ &= \frac{n}{z} \Pr[F_{n-1,1} > z] \\ &= \frac{n}{z} \left(1 - \left(\frac{z-1}{z}\right)^{n-1} \left(\frac{z-n}{z-1}\right) \right). \end{aligned}$$

Recall that for $z \leq n$, $\Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z] = 1$. To complete this proof, we use the formula $\mathbf{E}[\mathcal{F}^{(2)}(\mathbf{B}^{(n)})] = \int_0^\infty \Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z] dz = n + \int_n^\infty \Pr[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z] dz$.

In the form above, this is not easily integrable; however, we can transform it back into a binomial sum which we can integrate:

$$\begin{aligned}\Pr\left[\mathcal{F}^{(2)}(\mathbf{B}^{(n)}) > z\right] &= n \sum_{i=2}^n \left(\frac{-1}{z}\right)^i i \binom{n-1}{i-1}. \\ \mathbf{E}\left[\mathcal{F}^{(2)}(\mathbf{B}^{(n)})\right] &= n + n \int_n^\infty \sum_{i=2}^n \left(\frac{-1}{z}\right)^i i \binom{n-1}{i-1} dz \\ &= n - n \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}. \quad \blacksquare\end{aligned}$$

Theorem 2. *The competitive ratio of any auction on n bidders is*

$$1 - \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}.$$

This theorem comes from combining Lemma 2, Lemma 5, and Fact 1. Of course, for the special case of $n = 2$ this gives the lower bound of 2 that we already gave. For $n = 3$ this gives a lower bound of $13/6$. A lower bound for the competitive ratio of the best auction for general n is obtained by taking the limit. In the proof of the main theorem to follow, we use the following fact.

Fact 2 *For $1 \leq k \leq K$, $0 < a_k < 1$, then $\prod_{k=1}^K (1 - a_k) \geq 1 - \sum_{k=1}^K a_k$.*

Theorem 3. *The competitive ratio of any auction is at least 2.42.*

Proof. We prove this theorem by showing that,

$$\lim_{n \rightarrow \infty} \left(1 - \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}\right) = 1 + \sum_{i=2}^{\infty} (-1)^i \frac{i}{(i-1)(i-1)!} \quad (5)$$

After which, routine calculation shows that the right hand side of the above equation is at least 2.42 which gives the theorem. To prove that (5) holds, it is sufficient to show that

$$\left| \left(1 + \sum_{i=2}^n (-1)^i \frac{i}{(i-1)(i-1)!}\right) - \left(1 - \sum_{i=2}^n \left(\frac{-1}{n}\right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1}\right) \right| = O\left(\frac{1}{n}\right).$$

We proceed as follows:

$$\begin{aligned}
& \left| \left(1 + \sum_{i=2}^n (-1)^i \frac{i}{(i-1)(i-1)!} \right) - \left(1 - \sum_{i=2}^n \left(\frac{-1}{n} \right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1} \right) \right| \\
& \leq \sum_{i=2}^n \left| \frac{i}{(i-1)(i-1)!} - \left(\frac{1}{n} \right)^{i-1} \frac{i}{i-1} \binom{n-1}{i-1} \right| \\
& = \sum_{i=2}^n \left| \frac{i}{(i-1)(i-1)!} \left(1 - \frac{n(n-1) \cdots (n-i+2)}{n^{i-1}} \right) \right| \\
& = \sum_{i=2}^n \left| \frac{i}{(i-1)(i-1)!} \left(1 - \left(1 - \frac{1}{n} \right) \left(1 - \frac{2}{n} \right) \cdots \left(1 - \frac{i-2}{n} \right) \right) \right| \\
& \leq \sum_{i=2}^n \left| \frac{i}{(i-1)(i-1)!} \left(1 - \left(1 - \sum_{j=1}^{i-2} \frac{j}{n} \right) \right) \right| \\
& \leq \sum_{i=2}^n \left| \frac{i}{(i-1)(i-1)!} \left(\frac{i^2}{n} \right) \right| = \frac{1}{n} \sum_{i=2}^n \frac{i^3}{(i-1)(i-1)!} \leq \frac{1}{n} \sum_{i=2}^{\infty} \frac{i^3}{(i-1)(i-1)!}
\end{aligned}$$

Since $(i-1)!$ grows exponentially, $\sum_{i=2}^{\infty} \frac{i^3}{(i-1)(i-1)!}$ is bounded by a constant and we have the desired result. ■

4 Lower Bounds versus Upper Bounds

As mentioned earlier, the lower bound of 2.42 for large n does not match the competitive ratio of the best known auction (currently 3.39 [3]). In this section, we briefly consider the issue of matching upper bounds for small values of n . For $n = 2$ the 1-item Vickrey auction obtains the optimal competitive ratio of 2 (see Section 2.3). It is interesting to note that for the $n = 2$ case the optimal auction always uses sale prices chosen from the set of input bids (in particular, the second highest bid). This motivates the following definition.

Definition 6. *We say an auction, \mathcal{A} , is restricted if on any input the sale prices are drawn from the set of input bid values, unrestricted otherwise.*

While the Vickrey auction is a restricted auction, the Vickrey auction with reserve price r , which offers the highest bidder the greater of r and the second highest bid value, is not restricted as r may not necessarily be a bid value.

Designing restricted bid-independent auctions is easier than designing general ones as the set of sale prices is determined by the input bids. However, as we show next, even for the $n = 3$ case the optimal restricted auction's competitive ratio is worse than that of the optimal unrestricted auction.

Lemma 6. *For $n = 3$, no restricted truthful auction, BI_f , can achieve a competitive ratio better than $5/2$.*

Proof. Because BI_f is restricted, $f(a, b) \in \{a, b\}$. For $h > 1$ and $a \geq hb$, let

$$p = \sup_{a,b} \Pr[f(a, b) = b].$$

For ϵ close to zero, let a and b be such that $a > hb$ and $\Pr[f(a, b) = b] \geq p - \epsilon$.

The expected revenue for the auction on $\{a, b + \epsilon', b\}$ is at most $b + \epsilon' + pb$. Here, the $b + \epsilon'$ is an upper bound on the payment from the a bid and the pb is an upper bound on the expected revenue from the $b + \epsilon$ bid (as p is an upper bound on the probability that this bid is offered price b). Note that $\mathcal{F}^{(2)} = 3b$ so the competitive ratio obtained by taking the limit as $\epsilon' \rightarrow 0$ is at least $3/(1+p)$.

An upper bound for the expected revenue for the auction on $\{a + \epsilon', a, b\}$ is $2pb + (1-p+\epsilon)a$. The $2pb$ is from the $a + \epsilon'$ and the pb is from the a bid. For large h , $\mathcal{F}^{(2)} = 2a$ so the competitive ratio is at least $2h/(2pb + h(1-p+\epsilon))$. The limit as $\epsilon \rightarrow 0$ and $h \rightarrow \infty$ gives a bound on the competitive ratio of $2/(1-p)$.

Setting these two ratios equal we obtain an optimal value of $p = 1/5$ which obtains a competitive ratio of $5/2$. ■

This lower bound is tight as the following lemma shows.

Lemma 7. *For $a \geq b$, the bid-independent auction, BI_f with*

$$f(a, b) = \begin{cases} b & \text{with probability } 1/5 \\ a & \text{otherwise.} \end{cases}$$

achieves a competitive ratio of $5/2$ for three bidders.

We omit the proof as it follows via an elementary case analysis. It is interesting to note that the above auction is essentially performing a 1-item Vickrey auction with probability $4/5$ and a 2-item Vickrey auction with probability $1/5$.

Lemma 8. *An unrestricted three bidder auction can achieve a better competitive ratio than $5/2$.*

Proof. For $a \geq b$, the bid independent auction BI_f with

$$f(a, b) = \begin{cases} \begin{cases} b & \text{with probability } 15/23 \\ 3b/2 & \text{with probability } 8/23. \end{cases} & b \leq a \leq 3b/2 \\ \begin{cases} b & \text{with probability } 3/23 \\ a & \text{with probability } 20/23. \end{cases} & a > 3b/2. \end{cases}$$

has competitive ratio 2.3. We omit the elementary case analysis. ■

Recall that the lower bound on the competitive ratio for three bidders is $13/6 \approx 2.17$. Obtaining the optimal auction for three bidders remains an interesting open problem.

5 Conclusions

We have proven a lower bound of 2.42 on the competitive ratio of any truthful auction. The algorithmic technique used, that of looking at distributions of bidders on which all auctions perform the same and bounding the expected value of the metric (e.g., $\mathcal{F}^{(2)}$), is natural and useful for other auction related problems.

There is a strange artifact of the competitive framework that we employ here (and that which is used in prior work [2,3]). As we showed, the optimal worst case auction for selling two items is the 1-item Vickrey auction. This auction only sells one item, yet we had two items. Our optimal restricted auction for three items never sells more than two items. Yet, under our competitive framework it is not optimal to run this optimal restricted auction for three items when there are only two items for sale. As it turns out, this is not a problem when using a different but related metric, \mathcal{V}_{opt} , defined as the k -item Vickrey auction that obtains the highest profit, i.e., $\mathcal{V}_{\text{opt}}(\mathbf{b}) = \max_i (i-1)b_i$ (for $b_i \geq b_{i+1}$).

Acknowledgements

We would like to thank Amos Fiat for many helpful discussions.

References

1. Goldberg, A., Hartline, J., Wright, A.: Competitive Auctions and Digital Goods. In: Proc. 12th Symp. on Discrete Algorithms, ACM/SIAM (2001) 735–744
2. Fiat, A., Goldberg, A., Hartline, J., Karlin, A.: Competitive Generalized Auctions. In: Proc. 34th ACM Symposium on the Theory of Computing, ACM Press, New York (2002)
3. Goldberg, A., Hartline, J.: Competitiveness via Concensus. In: Proc. 14th Symp. on Discrete Algorithms, ACM/SIAM (2003)
4. Vickrey, W.: Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* **16** (1961) 8–37
5. Nisan, N., Ronen, A.: Algorithmic Mechanism Design. In: Proc. of 31st Symposium on Theory of Computing, ACM Press, New York (1999) 129–140
6. Archer, A., Tardos, E.: Truthful mechanisms for one-parameter agents. In: Proc. of the 42nd IEEE Symposium on Foundations of Computer Science. (2001)
7. Segal, I.: Optimal Pricing Mechanism with Unknown Demand. *American Economic Review* **93** (2003) 509–29
8. Lehmann, D., O’Callaghan, L., Shoham, Y.: Truth Revelation in Approximately Efficient Combinatorial Auctions. In: Proc. of 1st ACM Conf. on E-Commerce, ACM Press, New York (1999) 96–102
9. Mirrlees, J.: An Exploration into the Theory of Optimal Income Taxation. *Review of Economics Studies* **38** (1971) 175–208
10. Goldberg, A., Hartline, J., Karlin, A., Saks, M., Wright, A.: Competitive auctions and digital goods. *Games and Economic Behavior* (2002) Submitted for publication. An earlier version available as InterTrust Technical Report STAR-TR-99.09.01.
11. Abel, N.: Beweis eines Ausdrucks von welchem die Binomial-Formel ein einzelner Fall ist. *Crelles Journal für die Reine und Angewandte Mathematik* **1** (1826) 159–160