

From Optimal Limited To Unlimited Supply Auctions

Jason D. Hartline
Microsoft Research
1065 La Avenida)
Mountain View, CA 94043
hartline@microsoft.com

Robert McGrew^{*}
Computer Science Department
Stanford University
Stanford, CA 94305
bmcgrew@stanford.edu

ABSTRACT

We investigate the class of single-round, sealed-bid auctions for a set of identical items to bidders who each desire one unit. We adopt the worst-case competitive framework defined by [9, 5] that compares the profit of an auction to that of an optimal single-price sale of least two items. In this paper, we first derive an optimal auction for three items, answering an open question from [8]. Second, we show that the form of this auction is independent of the competitive framework used. Third, we propose a schema for converting a given limited-supply auction into an unlimited supply auction. Applying this technique to our optimal auction for three items, we achieve an auction with a competitive ratio of 3.25, which improves upon the previously best-known competitive ratio of 3.39 from [7]. Finally, we generalize a result from [8] and extend our understanding of the nature of the optimal competitive auction by showing that the optimal competitive auction occasionally offers prices that are higher than all bid values.

Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Non-numerical Algorithms and Problems; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Economics, Theory

Keywords

Auctions, Mechanism Design, Competitive Analysis

1. INTRODUCTION

The research area of *optimal mechanism design* looks at designing a mechanism to produce the most desirable outcome for the entity running the mechanism. This problem is well studied for the auction design problem where the optimal mechanism is the one

^{*}This work was supported in part by an NSF fellowship and NSF ITR 0205633.

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

EC'05, June 5–8, 2005, Vancouver, British Columbia, Canada.
Copyright 2005 ACM 1-59593-049-3/05/0006 ...\$5.00.

that brings the seller the most profit. Here, the classical approach is to design such a mechanism given the *prior distribution* from which the bidders' preferences are drawn (See e.g., [12, 4]). Recently Goldberg et al. [9] introduced the use of worst-case competitive analysis (See e.g., [3]) to analyze the performance of auctions that have no knowledge of the prior distribution. The goal of such work is to design an auction that achieves a large constant fraction of the profit attainable by an optimal mechanism that knows the prior distribution in advance. Positive results in this direction are fueled by the observation that in auctions for a number of identical units, much of the distribution from which the bidders are drawn can be deduced on the fly by the auction as it is being run [9, 14, 2].

The performance of an auction in such a worst-case competitive analysis is measured by its *competitive ratio*, the ratio between a benchmark performance and the auction's performance on the input distribution that maximizes this ratio. The holy grail of the worst-case competitive analysis of auctions is the auction that achieves the *optimal competitive ratio* (as small as possible). Since [9] this search has led to improved understanding of the nature of the optimal auction, the techniques for on-the-fly pricing in these scenarios, and the competitive ratio of the optimal auction [5, 7, 8]. In this paper we continue this line of research by improving in all of these directions. Furthermore, we give evidence corroborating the conjecture that the form of the optimal auction is independent of the benchmark used in the auction's competitive analysis. This result further validates the use of competitive analysis in gauging auction performance.

We consider the *single item, multi-unit, unit-demand* auction problem. In such an auction there are many units of a single item available for sale to bidders who each desire only one unit. Each bidder has a valuation representing how much the item is worth to him. 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 bid 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 falsely declare their valuations to get a better deal, we will adopt the solution concept of *truthful mechanism design* (see, e.g., [9, 15, 13]). In a truthful auction, revealing one's true valuation as one's 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.

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.

The competitive framework introduced in [9] and further refined in [5] 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 [9]. In this framework for competitive analysis, an auction is said to be β -competitive if it achieves a profit that is within a factor of $\beta \geq 1$ of the benchmark profit on every input. The optimal auction is the one which is β -competitive with the minimum value of β .

Previous to this work, the best known auction for the unlimited supply case had a competitive ratio of 3.39 [7] and the best lower bound known was 2.42 [8]. For the limited supply case, auctions can achieve substantially better competitive ratios. When there are only two units for sale, the optimal auction gives a competitive ratio of 2, which matches the lower bound for two units. When there are three units for sale, the best previously known auction had a competitive ratio of 2.3, compared with a lower bound of $13/6 \approx 2.17$ [8].

The results of this paper are as follows:

- We give the auction for three units that is optimally competitive against the profit of the omniscient single priced mechanism that sells at least two units. This auction achieves a competitive ratio of $13/6$, matching the lower bound from [8] (Section 3).
- We show that the form of the optimal auction is independent of the benchmark used in competitive analysis. In doing so, we give an optimal three bidder auction for generalized benchmarks (Section 4).
- We give a general technique for converting a limited supply auction into an unlimited supply auction where it is possible to use the competitive ratio of the limited supply auction to obtain a bound on the competitive ratio of the unlimited supply auction. We refer to auctions derived from this framework as *aggregation auctions* (Section 5).
- We improve on the best known competitive ratio by proving that the aggregation auction constructed from our optimal three-unit auction is 3.25-competitive (Section 5.1).
- Assuming that the conjecture that the optimal ℓ -unit auction has a competitive ratio that matches the lower bound proved in [8], we show that this optimal auction for $\ell \geq 3$ on some inputs will occasionally offer prices that are higher than any bid in that input (Section 6). For the three-unit case where we have shown that the lower bound of [8] is tight, this observation led to our construction of the optimal three-unit auction.

2. DEFINITIONS AND BACKGROUND

We consider single-round, sealed-bid auctions for a set of ℓ identical units of an item to bidders who each desire one unit. 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 [9]. Related formulations to this one have appeared in numerous places in recent literature (e.g., [1, 14, 5, 10]).

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. 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 [5].

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 terminology *bid-independent* and *truthful* interchangeably.

For a randomized bid-independent auction, $f(\mathbf{b}_{-i})$ is a random variable. We denote the probability density of $f(\mathbf{b}_{-i})$ at z by $\rho_{\mathbf{b}_{-i}}(z)$. We denote the profit of a truthful auction \mathcal{A} on input \mathbf{b} as $\mathcal{A}(\mathbf{b})$. The expected profit of the auction, $\mathbf{E}[\mathcal{A}(\mathbf{b})]$, is the sum of the expected payments made by each bidder, which we denote by $p_i(\mathbf{b})$ for bidder i . Clearly, the expected payment of each bid satisfies

$$p_i(\mathbf{b}) = \int_0^{b_i} x \rho_{\mathbf{b}_{-i}}(x) dx.$$

2.1 Competitive Framework

We now review the competitive framework from [5]. In order to evaluate the performance of auctions with respect to the goal of profit maximization, we introduce the optimal single price omniscient auction \mathcal{F} and the related omniscient auction $\mathcal{F}^{(2)}$.

DEFINITION 3. Give a vector $\mathbf{b} = (b_1, \dots, b_n)$, let $b_{(i)}$ represent the i -th largest value in \mathbf{b} .

The *optimal single price omniscient auction*, \mathcal{F} , is defined as follows. Auction \mathcal{F} on input \mathbf{b} determines the value k such that $kb_{(k)}$ is maximized. All bidders with $b_i \geq b_{(k)}$ win at price $b_{(k)}$; all remaining bidders lose. The profit of \mathcal{F} on input \mathbf{b} is thus $\mathcal{F}(\mathbf{b}) = \max_{1 \leq k \leq n} kb_{(k)}$.

In the competitive framework of [5] and subsequent papers, the performance of a truthful auction is gauged in comparison to $\mathcal{F}^{(2)}$, the *optimal singled priced auction that sells at least two units*. The profit of $\mathcal{F}^{(2)}$ is $\max_{2 \leq k \leq n} kb_{(k)}$. There are a number of reasons to choose this benchmark for comparison, interested readers should see [5] or [6] for a more detailed discussion.

Let \mathcal{A} be a truthful auction. We say that \mathcal{A} is β -competitive against $\mathcal{F}^{(2)}$ (or just β -competitive) if for all bid vectors \mathbf{b} , the expected profit of \mathcal{A} on \mathbf{b} satisfies

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

In Section 4 we generalize this framework to other profit benchmarks.

2.2 Scale Invariant and Symmetric Auctions

A *symmetric* auction is one where the auction outcome is unchanged when the input bids arrive in a different permutation. Goldberg et al. [8] show that a symmetric auction achieves the optimal competitive ratio. This is natural as the profit benchmark we consider is symmetric, and it allows us to consider only symmetric auctions when looking for the one with the optimal competitive ratio.

An auction defined by bid-independent function f is *scale invariant* if, for all i and all z , $\Pr[f(\mathbf{b}_{-i}) \geq z] = \Pr[f(c\mathbf{b}_{-i}) \geq cz]$. It is conjectured that the assumption of scale invariance is without loss of generality. Thus, we are motivated to consider symmetric scale-invariant auctions. When specifying a symmetric scale-invariant auction we can assume that f is only a function of the relative magnitudes of the $n - 1$ bids in \mathbf{b}_{-i} and that one of the bids, $b_j = 1$. It will be convenient to specify such auctions via the density function of $f(\mathbf{b}_{-i})$, $\rho_{\mathbf{b}_{-i}}(z)$. It is enough to specify such a density function of the form $\rho_{1,z_1,\dots,z_{n-1}}(z)$ with $1 \leq z_i \leq z_{i+1}$.

2.3 Limited Supply Versus Unlimited Supply

Following [8], 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 units for sale. This is without loss of generality as (a) any lower bound that applies to the $n = \ell$ case also extends to the case where $n \geq \ell$ [8], and (b) there is a reduction from the unlimited supply auction problem to the limited supply auction problem that takes an unlimited supply auction that is β -competitive with $\mathcal{F}^{(2)}$ and constructs a limited supply auction parameterized by ℓ that is β -competitive with $\mathcal{F}^{(2,\ell)}$, the optimal omniscient auction that sells between 2 and ℓ units [6].

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.

2.4 Lower Bounds and Optimal Auctions

Frequently in this paper, we will refer to the best known lower bound on the competitive ratio of truthful auctions:

THEOREM 2. [8] *The competitive ratio of any auction on n bidders is at least*

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

DEFINITION 4. *Let Υ_n denote the n -bidder auction that achieves the optimal competitive ratio.*

This bound is derived by analyzing the performance of any auction on the following distribution \mathbf{B} . In each random bid vector \mathbf{B} , each bid B_i is drawn i.i.d. from the distribution such that $\Pr[B_i \geq s] \leq 1/s$ for all $s \in \mathcal{S}$.

In the two-bidder case, this lower bound is 2. This is achieved by Υ_2 which is the 1-unit Vickrey auction.¹ In the three-bidder case, this lower bound is 13/6. In the next section, we define the auction Υ_3 which matches this lower bound. In the four-bidder case, this lower bound is 96/215. In the limit as the number of bidders grows, this lower bound approaches a number which is approximately 2.42.

It is conjectured that this lower bound is tight for any number of bidders and the optimal auction, Υ_n , matches it.

¹The 1-unit Vickrey auction sells to the highest bidder at the second highest bid value.

2.5 Profit Extraction

In this section we review the truthful profit extraction mechanism ProfitExtract_R . This mechanism is a special case of a general cost-sharing schema due to Moulin and Shenker [11].

The goal of profit extraction is, given bids \mathbf{b} , to extract a target value R of profit from some subset of the bidders.

ProfitExtract_R: Given bids \mathbf{b} , find the largest k such that the highest k bidders can equally share the cost R . Charge each of these bidders R/k . If no subset of bidders can cover the cost, the mechanism has no winners.

Important properties of this auction are as follows:

- ProfitExtract_R is truthful.
- If $R \leq \mathcal{F}(\mathbf{b})$, $\text{ProfitExtract}_R(\mathbf{b}) = R$; otherwise it has no winners and no revenue.

We will use this profit extraction mechanism in Section 5 with the following intuition. Such a profit extractor makes it possible to treat this subset of bidders as a single bid with value $\mathcal{F}(S)$. Note that given a single bid, b , a truthful mechanism might offer it price t and if $t \leq b$ then the bidder wins and pays t ; otherwise the bidder pays nothing (and loses). Likewise, a mechanism can offer the set of bidders S a target revenue R . If $R \leq \mathcal{F}^{(2)}(S)$, then ProfitExtract_R raises R from S ; otherwise, it raises no revenue from S .

3. AN OPTIMAL AUCTION FOR THREE BIDDERS

In this section we define the optimal auction for three bidders, Υ_3 , and prove that it indeed matches the known lower bound of 13/6. We follow the definition and proof with a discussion of how this auction was derived.

DEFINITION 5. *Υ_3 is scale-invariant and symmetric and given by the bid-independent function with density function*

$$\rho_{1,x}(z) = \begin{cases} \text{For } x \leq 3/2 \\ \left\{ \begin{array}{l} 1 \text{ with probability } 9/13 \\ z \text{ with probability density } g(z) \text{ for } z > 3/2 \end{array} \right. \\ \text{For } x > 3/2 \\ \left\{ \begin{array}{l} 1 \text{ with probability } 9/13 - \int_{3/2}^x z g(z) dz \\ x \text{ with probability } \int_{3/2}^x (z+1) g(z) dz \\ z \text{ with probability density } g(z) \text{ for } z > x \end{array} \right. \end{cases}$$

where $g(x) = \frac{2/13}{(x-1)^3}$.

THEOREM 3. *The Υ_3 auction has a competitive ratio of 13/6 \approx 2.17, which is optimal. Furthermore, the auction raises exactly $\frac{6}{13} \mathcal{F}^{(2)}$ on every input with non-identical bids.*

PROOF. Consider the bids $1, x, y$, with $1 < x < y$. There are three cases.

CASE 1 ($x < y \leq 3/2$): $\mathcal{F}^{(2)} = 3$. The auction must raise expected revenue of at least 18/13 on these bids. The bidder with valuation x will pay 1 with 9/13, and the bidder with valuation y will pay 1 with probability 9/13. Therefore Υ_3 raises 18/13 on these bids.

CASE 2 ($x \leq 3/2 < y$): $\mathcal{F}^{(2)} = 3$. The auction must raise expected revenue of at least 18/13 on these bids. The bidder with

valuation x will pay $9/13 - \int_{3/2}^y z g(z) dz$ in expectation. The bidder with valuation y will pay $9/13 + \int_{3/2}^y z g(z) dz$ in expectation. Therefore Υ_3 raises $18/13$ on these bids.

CASE 3 ($3/2 < x \leq y$): $\mathcal{F}^{(2)} = 2x$. The auction must raise expected revenue of at least $12x/13$ on these bids. Consider the revenue raised from all three bidders:

$$\begin{aligned} \mathbb{E}[\Upsilon_3(\mathbf{b})] &= p(1, x, y) + p(x, 1, y) + p(y, 1, x) \\ &= 0 + 9/13 - \int_{3/2}^y z g(z) dz + 9/13 - \int_{3/2}^x z g(z) dz \\ &\quad + x \int_{3/2}^x (z+1) g(z) dz + \int_x^y z g(z) dz \\ &= 18/13 + (x-2) \int_{3/2}^x z g(z) dz + x \int_{3/2}^x g(z) dz \\ &= 12x/13. \end{aligned}$$

The final equation comes from substituting in $g(x) = \frac{2/13}{(x-1)^3}$ and expanding the integrals. Note that the fraction of $\mathcal{F}^{(2)}$ raised on every input is identical. If any of the inequalities $1 \leq x \leq y$ are not strict, the same proof applies giving a lower bound on the auction's profit; however, this bound may no longer be tight. \square

Motivation for Υ_3

In this section, we will conjecture that a particular input distribution is worst-case, and show, as a consequence, that all inputs are worst-case in the optimal auction. By applying this consequence, we will derive an optimal auction for three bidders.

A truthful, randomized auction on n bidders can be represented by a randomized function $f : \mathbb{R}^{n-1} \times n \rightarrow \mathbb{R}$ that maps masked bid vectors to prices in \mathbb{R} . By normalization, we can assume that the lowest possible bid is 1. Recall that $\rho_{\mathbf{b}_i}(z) = \Pr[f(\mathbf{b}_i) = z]$.

The optimal auction for the finite auction problem can be found by the following optimization problem in which the variables are $\rho_{\mathbf{b}_i}(z)$:

$$\begin{aligned} &\text{maximize} && r \\ &\text{subject to} && \sum_{i=1}^n \int_{z=1}^{b_i} z \rho_{\mathbf{b}_i}(z) \geq r \mathcal{F}^{(2)}(\mathbf{b}) \\ &&& \int_{z=1}^{\infty} \rho_{\mathbf{b}_i}(z) = 1 \\ &&& \rho_{\mathbf{b}_i}(z) \geq 0 \end{aligned}$$

This set of integral inequalities is difficult to maximize over. However, by guessing which constraints are tight and which are slack at the optimum, we will be able to derive a set of differential equations for which any feasible solution is an optimal auction.

As we discuss in Section 2.4, in [8], the authors define a distribution and use it to find a lower bound on the competitive ratio of the optimal auction. For two bidders, this bid distribution is the worst-case input distribution. We guess (and later verify) that this distribution is the worst-case input distribution for three bidders as well. Since this distribution has full support over the set of all bid vectors and a worst-case distribution puts positive probability only on worst-case inputs, we can therefore assume that all but a measure zero set of inputs is worst-case for the optimal auction. In the optimal two-bidder auction, all inputs with non-identical bids are worst-case, so we will assume the same for three bidders.

The guess that these constraints are tight allows us to transform the optimization problem into a feasibility problem constrained by differential equations. If the solution to these equations has value matching the lower bound obtained from the worst-case distribu-

tion, then this solution is the optimal auction and that our conjectured choice of worst-case distribution is correct.

In Section 6 we show that the optimal auction must sometimes place probability mass on sale prices above the highest bid. This motivates considering symmetric scale-invariant auctions for three bidders with probability density function, $\rho_{1,x}(z)$, of the following form:

$$\rho_{1,x}(z) = \begin{cases} 1 & \text{with discrete probability } a(x) \\ x & \text{with discrete probability } b(x) \\ z & \text{with probability density } g(z) \text{ for } z > x \end{cases}$$

In this auction, the sale price for the first bidder is either one of the latter two bids, or higher than either bid with a probability density which is independent of the input.

The feasibility problem which arises from the linear optimization problem by assuming the constraints are tight is as follows:

$$\begin{aligned} a(y) + a(x) + xb(x) + \int_x^y z g(z) dz &= r \max(3, 2x) \quad \forall x < y \\ a(x) + b(x) + \int_x^{\infty} g(z) dz &= 1 \\ a(x) &\geq 0 \\ b(x) &\geq 0 \\ g(z) &\geq 0 \end{aligned}$$

Solving this feasibility problem gives the auction Υ_3 proposed above. The proof of its optimality validates its proposed form. Finding a simple restriction on the form of n -bidder auctions for $n > 3$ under which the optimal auction can be found analytically as above remains an open problem.

4. GENERALIZED PROFIT BENCHMARKS

In this section, we widen our focus beyond auctions that compete with $\mathcal{F}^{(2)}$ to consider other benchmarks for an auction's profit. We will show that, for three bidders, the form of the optimal auction is essentially independent of the benchmark profit used. This results strongly corroborates the worst-case competitive analysis of auctions by showing that our techniques allow us to derive auctions which are competitive against a broad variety of reasonable benchmarks rather than simply against $\mathcal{F}^{(2)}$.

Previous work in competitive analysis of auctions has focused on the question of designing the auction with the best competitive ratio against $\mathcal{F}^{(2)}$, the profit of the optimal omniscient single-priced mechanism that sells at least two items. However, it is reasonable to consider other benchmarks. For instance, one might wish to compete against \mathcal{V}^* , the profit of the k -Vickrey auction with optimal-in-hindsight choice of k .² Alternatively, if an auction is being used as a subroutine in a larger mechanism, one might wish to choose the auction which is optimally competitive with a benchmark specific to that purpose.

Recall that $\mathcal{F}^{(2)}(\mathbf{b}) = \max_{2 \leq k \leq n} kb_{(k)}$. We can generalize this definition to \mathcal{G}_s , parameterized by $\mathbf{s} = (s_2, \dots, s_n)$ and defined as:

$$\mathcal{G}_s(\mathbf{b}) = \max_{2 \leq k \leq n} s_k b_{(k)}.$$

When considering \mathcal{G}_s we assume without loss of generality that $s_i < s_{i+1}$ as otherwise the constraint imposed by s_{i+1} is irrelevant. Note that $\mathcal{F}^{(2)}$ is the special case of \mathcal{G}_s with $s_i = i$, and that $\mathcal{V}^* = \mathcal{G}_s$ with $s_i = i - 1$.

²Recall that the k -Vickrey auction sells a unit to each of the highest k bidders at a price equal to the $k + 1$ st highest bid, $b_{(k+1)}$, achieving a profit of $kb_{(k+1)}$.

Competing with \mathcal{G}_s

We will now design a three-bidder auction $\Upsilon_3^{s,t}$ that achieves the optimal competitive ratio against $\mathcal{G}_{s,t}$. As before, we will first find a lower bound on the competitive ratio and then design an auction to meet that bound.

We can lower bound the competitive ratio of $\Upsilon_3^{s,t}$ using the same worst-case distribution from [8] that we used against $\mathcal{F}^{(2)}$. Evaluating the performance of any auction competing against $\mathcal{G}_{s,t}$ on this distribution will yield the following theorem. We denote the optimal auction for three bidders against $\mathcal{G}_{s,t}$ as $\Upsilon_3^{s,t}$.

THEOREM 4. *The optimal three-bidder auction, $\Upsilon_3^{s,t}$, competing against $\mathcal{G}_{s,t}(\mathbf{b}) = \max(sb_{(2)}, tb_{(3)})$ has a competitive ratio of at least $\frac{s^2+t^2}{2t}$.*

The proof can be found in the appendix.

Similarly, we can find the optimal auction against $\mathcal{G}_{s,t}$ using the same technique we used to solve for the three bidder auction with the best competitive ratio against $\mathcal{F}^{(2)}$.

DEFINITION 6. $\Upsilon_3^{s,t}$ is scale-invariant and symmetric and given by the bid-independent function with density function

$$\rho_{1,x}(z) = \begin{cases} \text{For } x \leq \frac{t}{s} \\ \left\{ \begin{array}{l} 1 \text{ with probability } \frac{t^2}{s^2+t^2} \\ z \text{ with probability density } g(z) \text{ for } z > \frac{t}{s} \end{array} \right. \\ \text{For } x > \frac{t}{s} \\ \left\{ \begin{array}{l} 1 \text{ with probability } \frac{t^2}{s^2+t^2} - \int_{\frac{t}{s}}^x zg(z)dz \\ x \text{ with probability } \int_{\frac{t}{s}}^x (z+1)g(z)dz \\ z \text{ with probability density } g(z) \text{ for } z > x \end{array} \right. \end{cases}$$

where $g(x) = \frac{2(t-s)^2/(s^2+t^2)}{(x-1)^3}$.

THEOREM 5. $\Upsilon_3^{s,t}$ is $\frac{s^2+t^2}{2t}$ -competitive with $\mathcal{G}_{s,t}$.

This auction, like Υ_3 , can be derived by reducing the optimization problem to a feasibility problem, guessing that the optimal solution has the same form as $\Upsilon_3^{s,t}$, and solving. The auction is optimal because it matches the lower bound found above. Note that the form of $\Upsilon_3^{s,t}$ is essentially the same as for Υ_3 , but that the probability of each price is scaled depending on the values of s and t .

That our auction for three bidders matches the lower bound computed by the input distribution used in [8] is strong evidence that this input distribution is the worst-case input distribution for any number of bidders and any generalized profit benchmark. Furthermore, we strongly suspect that for any number of bidders, the form of the optimal auction will be independent of the benchmark used.

5. AGGREGATION AUCTIONS

We have seen that optimal auctions for small cases of the limited-supply model can be found analytically. In this section, we will construct a schema for turning limited supply auctions into unlimited supply auctions with a good competitive ratio.

As discussed in Section 2.5, the existence of a profit extractor, ProfitExtract_R , allows an auction to treat a set of bids S as a single bid with value $\mathcal{F}(S)$. Given n bidders and an auction, \mathcal{A}_m , for $m < n$ bidders, we can convert the m -bidder auction into an n -bidder auction by randomly partitioning the bidders into m subsets and then treating each subset as a single bidder (via ProfitExtract_R) and running the m -bidder auction.

DEFINITION 7. *Given a truthful m -bidder auction, \mathcal{A}_m , the m -aggregation auction for \mathcal{A}_m , $\text{Agg}_{\mathcal{A}_m}$, works as follows:*

1. Cast each bid uniformly at random into one of m bins, resulting in bid vectors $\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(m)}$.
2. For each bin j , compute the aggregate bid $B_j = \mathcal{F}(\mathbf{b}^{(j)})$. Let \mathbf{B} be the vector of aggregate bids, and \mathbf{B}_{-j} be the aggregate bids for all bins other than j .
3. Compute the aggregate price $T_j = f(\mathbf{B}_{-j})$, where f is the bid-independent function for \mathcal{A}_m .
4. For each bin j , run $\text{ProfitExtract}_{T_j}$ on $\mathbf{b}^{(j)}$.

Since \mathcal{A}_m and ProfitExtract_R are truthful, T_j is computed independently of any bid in bin j and thus the price offered any bidder in $\mathbf{b}^{(j)}$ is independent of his bid; therefore,

THEOREM 6. *If \mathcal{A}_m is truthful, the m -aggregation auction for \mathcal{A}_m , $\text{Agg}_{\mathcal{A}_m}$, is truthful.*

Note that this schema yields a new way of understanding the Random Sampling Profit Extraction (RSPE) auction [5] as the simplest case of an aggregation auction. It is the 2-aggregation auction for Υ_2 , the 1-unit Vickrey auction.

To analyze $\text{Agg}_{\mathcal{A}_m}$, consider throwing k balls into m labeled bins. Let \mathbf{k} represent a configuration of balls in bins, so that k_i is equal to the number of balls in bin i , and $k_{(i)}$ is equal to the number of balls in the i th largest bin. Let $\mathbf{K}_{m,k}$ represent the set of all possible configurations of k balls in m bins. We write the multinomial coefficient of \mathbf{k} as $\binom{k}{\mathbf{k}}$. The probability that a particular configuration \mathbf{k} arises by throwing balls into bins uniformly at random is $\binom{k}{\mathbf{k}} m^{-k}$.

THEOREM 7. *Let \mathcal{A}_m be an auction with competitive ratio β . Then the m -aggregation auction for \mathcal{A}_m , $\text{Agg}_{\mathcal{A}_m}$, raises the following fraction of the optimal revenue $\mathcal{F}^{(2)}(\mathbf{b})$:*

$$\frac{\mathbf{E}[\text{Agg}_{\mathcal{A}_m}(\mathbf{b})]}{\mathcal{F}^{(2)}} \geq \min_{k \geq 2} \sum_{\mathbf{k} \in \mathbf{K}_{m,k}} \frac{\mathcal{F}^{(2)}(\mathbf{k}) \binom{k}{\mathbf{k}}}{\beta k m^k}$$

PROOF. By definition, $\mathcal{F}^{(2)}$ sells to $k \geq 2$ bidders at a single price p . Let k_j be the number of such bidders in $\mathbf{b}^{(j)}$. Clearly, $\mathcal{F}(\mathbf{b}^{(j)}) \geq pk_j$. Therefore,

$$\begin{aligned} \frac{\mathcal{F}^{(2)}(\mathcal{F}(\mathbf{b}^{(1)}), \dots, \mathcal{F}(\mathbf{b}^{(n)}))}{\mathcal{F}^{(2)}(\mathbf{b})} &\geq \frac{\mathcal{F}^{(2)}(pk_1, \dots, pk_n)}{pk} \\ &= \frac{\mathcal{F}^{(2)}(k_1, \dots, k_n)}{k} \end{aligned}$$

The inequality follows from the monotonicity of $\mathcal{F}^{(2)}$, and the equality from the homogeneity of $\mathcal{F}^{(2)}$.

$\text{ProfitExtract}_{T_j}$ will raise T_j if $T_j \leq B_j$, and no profit otherwise. Thus, $\mathbf{E}[\text{Agg}_{\mathcal{A}_m}(\mathbf{b})] \geq \mathbf{E}[\mathcal{F}^{(2)}(\mathbf{B})/\beta]$. The theorem follows by rewriting this expectation as a sum over all \mathbf{k} in $\mathbf{K}_{m,k}$. \square

5.1 A 3.25 Competitive Auction

We apply the aggregation auction schema to Υ_3 , our optimal auction for three bidders, to achieve an auction with competitive ratio 3.25. This improves on the previously best known auction which is 3.39-competitive [7].

THEOREM 8. *The aggregation auction for Υ_3 has competitive ratio 3.25.*

PROOF. By theorem 7,

$$\frac{\mathbf{E}[\text{Agg}_{\Upsilon_3}(\mathbf{b})]}{\mathcal{F}^{(2)}(\mathbf{b})} \geq \min_{k \geq 2} \sum_{i=1}^k \sum_{j=1}^{k-i} \frac{\mathcal{F}^{(2)}(i, j, k-i-j) \binom{k}{i, j, k-i-j}}{\beta k 3^k}$$

For $k = 2$ and $k = 3$, $\mathbf{E}[\text{Agg}_{\Upsilon_3}(\mathbf{b})] = \frac{2}{3}k/\beta$. As k increases, $\mathbf{E}[\text{Agg}_{\Upsilon_3}(\mathbf{b})]/\mathcal{F}^{(2)}$ increases as well. Since we do not expect to find a closed-form formula for the revenue, we lower bound $\mathcal{F}^{(2)}(\mathbf{b})$ by $3b_{(3)}$. Using large deviation bounds, one can show that this lower bound is greater than $\frac{2}{3}k/\beta$ for large-enough k , and the remainder can be shown by explicit calculation.

Plugging in $\beta = 13/6$, the competitive ratio is $13/4$. As k increases, the competitive ratio approaches $13/6$.

Note that the above bound on the competitive ratio of Agg_{Υ_3} is tight. To see this, consider the bid vector with two very large and non-identical bids of h and $h + \epsilon$ with the remaining bids 1. Given that the competitive ratio of Υ_3 is tight on this example, the expected revenue of this auction on this input will be exactly $13/4$. \square

5.2 A $\mathcal{G}_{s,t}$ -based Aggregation Auction

In this section we show that Υ_3 is not the optimal auction to use in an aggregation auction. One can do better by choosing the auction that is optimally competitive against a specially tailored benchmark.

To see why this might be the case, notice (Table 1) that the fraction of $\mathcal{F}^{(2)}(\mathbf{b})$ raised for when there are $k = 2$ and $k = 3$ winning bidders in $\mathcal{F}^{(2)}(\mathbf{b})$ is substantially smaller than the fraction of $\mathcal{F}^{(2)}(\mathbf{b})$ raised when there are more winners. This occurs because the expected ratio between $\mathcal{F}^{(2)}(\mathbf{B})$ and $\mathcal{F}^{(2)}(\mathbf{b})$ is lower in this case while the competitive ratio of Υ_3 is constant. If we chose a three bidder auction that performed better when $\mathcal{F}^{(2)}$ has smaller numbers of winners, our aggregation auction would perform better in the worst case.

One approach is to compete against a different benchmark that puts more weight than $\mathcal{F}^{(2)}$ on solutions with a small number of winners. Recall that $\mathcal{F}^{(2)}$ is the instance of $\mathcal{G}_{s,t}$ with $s = 2$ and $t = 3$. By using the auction that competes optimally against $\mathcal{G}_{s,t}$ with $s > 2$, while holding $t = 3$, we will raise a higher fraction of revenue on smaller numbers of winning bidders and a lower fraction of revenue on large numbers of winning bidders. We can numerically optimize the values of s and t in $\mathcal{G}_{s,t}(\mathbf{b})$ in order to achieve the best competitive ratio for the aggregation auction. In fact, this will allow us to improve our competitive ratio slightly.

THEOREM 9. *For an optimal choice of s and t , the aggregation auction for $\Upsilon_3^{s,t}$ is 3.243-competitive.*

The proof follows the outline of Theorem 7 and 8 with the optimal choice of $s = 2.162$ (while t is held constant at 3).

5.3 Further Reducing the Competitive Ratio

There are a number of ways we might attempt to use this aggregation auction schema to continue to push the competitive ratio down. In this section, we give a brief discussion of several attempts.

5.3.1 Agg_{Υ_m} for $m > 3$

If the aggregation auction for Υ_2 has a competitive ratio of 4 and the aggregation auction for Υ_3 has a competitive ratio of 3.25, can we improve the competitive ratio by aggregating Υ_4 or Υ_m for larger m ? We conjecture in the negative: for $m > 3$, the aggregation auction for Υ_m has a larger competitive ratio than the aggregation auction for Υ_3 . The primary difficulty in proving this

k	m = 2	m = 3	m = 4	m = 5	m = 6	m = 7
2	0.25	0.3077	0.3349	0.3508	0.3612	0.3686
3	0.25	0.3077	0.3349	0.3508	0.3612	0.3686
4	0.3125	0.3248	0.3349	0.3438	0.3512	0.3573
5	0.3125	0.3191	0.3244	0.3311	0.3378	0.3439
6	0.3438	0.321	0.3057	0.3056	0.311	0.318
7	0.3438	0.333	0.3081	0.3009	0.3025	0.3074
8	0.3633	0.3229	0.3109	0.3022	0.3002	0.3024
9	0.3633	0.3233	0.3057	0.2977	0.2927	0.292
10	0.377	0.3328	0.308	0.2952	0.2866	0.2837
11	0.377	0.3319	0.3128	0.298	0.2865	0.2813
12	0.3872	0.3358	0.3105	0.3001	0.2894	0.2827
13	0.3872	0.3395	0.3092	0.2976	0.2905	0.2841
14	0.3953	0.3391	0.312	0.2961	0.2888	0.2835
15	0.3953	0.3427	0.3135	0.2973	0.2882	0.2825
16	0.4018	0.3433	0.3128	0.298	0.2884	0.2823
17	0.4018	0.3428	0.3129	0.2967	0.2878	0.282
18	0.4073	0.3461	0.3133	0.2959	0.2859	0.2808
19	0.4073	0.3477	0.3137	0.2962	0.2844	0.2789
20	0.4119	0.3486	0.3148	0.2973	0.2843	0.2777
21	0.4119	0.3506	0.3171	0.298	0.2851	0.2775
22	0.4159	0.3519	0.3189	0.2986	0.2863	0.2781
23	0.4159	0.3531	0.3202	0.2995	0.2872	0.2791
24	0.4194	0.3539	0.3209	0.3003	0.2878	0.2797
25	0.4194	0.3548	0.3218	0.3012	0.2886	0.2801

Table 1: $\mathbf{E}[\mathcal{A}(\mathbf{b})/\mathcal{F}^{(2)}(\mathbf{b})]$ for Agg_{Υ_m} as a function of k , the optimal number of winners in $\mathcal{F}^{(2)}(\mathbf{b})$. The lowest value for each column is printed in bold.

conjecture lies in the difficulty of finding a closed-form solution for the formula of Theorem 7. We can, however, evaluate this formula numerically for different values of m and k , assuming that the competitive ratio for Υ_m matches the lower bound for m given by Theorem 2. Table 1 shows, for each value of m and k , the fraction of $\mathcal{F}^{(2)}$ raised by the aggregation auction for Agg_{Υ_m} when there are k winning bidders, assuming the lower bound of Theorem 2 is tight.

5.3.2 Convex combinations of Agg_{Υ_m}

As can be seen in Table 1, when $m > 3$, the worst-case value of k is no longer 2 and 3, but instead an increasing function of m . An aggregation auction for Υ_m outperforms the aggregation auction for Υ_3 when there are two or three winning bidders, while the aggregation auction for Υ_3 outperforms the other when there are at least six winning bidders. Thus, for instance, an auction which randomizes between aggregation auctions for Υ_3 and Υ_4 will have a worst-case which is better than that of either auction alone. Larger combinations of auctions will allow more room to optimize the worst-case. However, we suspect that no convex combination of aggregation auctions will have a competitive ratio lower than 3. Furthermore, note that we cannot yet claim the existence of a good auction via this technique as the optimal auction Υ_n for $n > 3$ is not known and it is only conjectured that the bound given by Theorem 2 and represented in Table 1 is correct for Υ_n .

6. A LOWER BOUND FOR CONSERVATIVE AUCTIONS

In this section, we define a class of auctions that never offer a sale price which is higher than any bid in the input and prove a lower bound on the competitive ratio of these auctions. As this

lower bound is stronger than the lower bound of Theorem 2 for $n \geq 3$, it shows that the optimal auction must occasionally charge a sales price higher than any bid in the input. Specifically, this result partially explains the form of the optimal three bidder auction.

DEFINITION 8. We say an auction \mathcal{A} is conservative if its bid-independent function f satisfies $f(\mathbf{b}_{-i}) \leq \max(\mathbf{b}_{-i})$.

We can now state our lower bound for conservative auctions.

THEOREM 10. Let \mathcal{A} be a conservative auction for n bidders. Then the competitive ratio of \mathcal{A} is at least $\frac{3n-2}{n}$.

COROLLARY 1. The competitive ratio of any conservative auction for an arbitrary number of bidders is at least three.

For a two-bidder auction, this restriction does not prevent optimality. Υ_2 , the 1-unit Vickrey auction, is conservative. For larger numbers of bidders, however, the restriction to conservative auctions does affect the competitive ratio. For the three-bidder case, Υ_3 has competitive ratio 2.17, while the best conservative auction is no better than 2.33-competitive.

The k -Vickrey auction and the Random Sampling Optimal Price auction [9] are conservative auctions. The Random Sampling Profit Extraction auction [5] and the CORE auction [7], on the other hand, use the ProfitExtract_R mechanism as a subroutine and thus sometimes offer a sale price which is higher than the highest input bid value.

In [8], the authors define a *restricted* auction as one on which, for any input, the sale prices are drawn from the set of input bid values. The class of conservative auctions can be viewed as a generalization of the class of restricted auctions and therefore our result below gives lower bounds on the performance of a broader class of auctions.

We will prove Theorem 10 with the aid of the following lemma:

LEMMA 1. Let \mathcal{A} be a conservative auction with competitive ratio $1/r$ for n bidders. Let $h \gg n$. Let $h_0 = 1$ and $h_k = kh$ otherwise. Then, for all k and $H \geq kh$, $\Pr[f(1, 1, \dots, 1, H) \leq h_k] \geq \frac{nr-1}{n-1} + k(\frac{3nr-2r-n}{n-1})$.

PROOF. The lemma is proved by strong induction on k . First some notation that will be convenient. For any particular k and H we will be considering the bid vector $\mathbf{b} = (1, \dots, 1, h_k, H)$ and placing bounds on $\rho_{\mathbf{b}_i}(z)$. Since we can assume without loss of generality that the auction is symmetric, we will notate \mathbf{b}_{-1} as \mathbf{b} with any one of the 1-valued bids masked. Similarly we notate \mathbf{b}_{-h_k} (resp. \mathbf{b}_{-H}) as \mathbf{b} with the h_k -valued bid (resp. H -valued bid) masked. We will also let $p_1(\mathbf{b})$, $p_{h_k}(\mathbf{b})$, and $p_H(\mathbf{b})$ represent the expected payment of a 1-valued, h_k -valued, and H -valued bidder in \mathcal{A} on \mathbf{b} , respectively (note by symmetry the expected payment for all 1-valued bidders is the same).

Base case ($k = 0, h_k = 1$): \mathcal{A} must raise revenue of at least rn on $\mathbf{b} = (1, \dots, 1, 1, H)$:

$$\begin{aligned} rn &\leq p_H(\mathbf{b}) + (n-1)p_1(\mathbf{b}) \\ &= 1 + (n-1) \int_0^1 x \rho_{\mathbf{b}_{-1}}(x) dx \\ &\leq 1 + (n-1) \int_0^1 \rho_{\mathbf{b}_{-1}}(x) dx \end{aligned}$$

The second inequality follows from the conservatism of the underlying auction. The base case follows trivially from the final inequality.

Inductive case ($k > 0, h_k = kh$): Let $\mathbf{b} = (1, \dots, 1, h_k, H)$. First, we will find an upper bound on $p_H(\mathbf{b})$

$$\begin{aligned} p_H(\mathbf{b}) &= \int_0^1 x \rho_{\mathbf{b}_{-H}}(x) dx + \sum_{i=1}^k \int_{h_{i-1}}^{h_i} x \rho_{\mathbf{b}_{-H}}(x) dx \quad (1) \\ &\leq 1 + \sum_{i=1}^k h_i \int_{h_{i-1}}^{h_i} \rho_{\mathbf{b}_{-H}}(x) dx \\ &\leq 1 + \left(\frac{3nr-2r-n}{n-1} \right) \sum_{i=1}^{k-1} ih \\ &\quad + kh \left(1 - \frac{nr-1}{n-1} - (k-1) \frac{3nr-2r-n}{n-1} \right) \quad (2) \\ &= kh \left[\frac{n(1-r)}{n-1} + \frac{(k-1)3nr-2r-n}{2(n-1)} \right] + 1. \quad (3) \end{aligned}$$

Equation (1) follows from the conservatism of \mathcal{A} and (2) is from invoking the strong inductive hypothesis with $H = kh$ and the observation that the maximum possible revenue will be found by placing exactly enough probability at each multiple of h to satisfy the constraints of the inductive hypothesis and placing the remaining probability at kh . Next, we will find a lower bound on $p_{h_k}(\mathbf{b})$ by considering the revenue raised by the bids \mathbf{b} . Recall that \mathcal{A} must obtain a profit of at least $r\mathcal{F}^{(2)}(\mathbf{b}) = 2rkh$. Given upper-bounds on the profit from the H -valued, equation bid (3), and the 1-valued bids, the profit from the h_k -valued bid must be at least:

$$\begin{aligned} p_{h_k}(\mathbf{b}) &\geq 2rkh - (n-2)p_1(\mathbf{b}) - p_H(\mathbf{b}) \\ &\geq kh \left[2r - \frac{n(1-r)}{n-1} + \frac{(k-1)3nr-2r-n}{2(n-1)} \right] - O(n). \quad (4) \end{aligned}$$

In order to lower bound $\Pr[f(\mathbf{b}_{-h_k}) \leq kh]$, consider the auction that minimizes it and is consistent with the lower bounds obtained by the strong inductive hypothesis on $\Pr[f(\mathbf{b}_{-h_k}) \leq ih]$. To minimize the constraints implied by the strong inductive hypothesis, we place the minimal amount of probability mass required each price level. This gives $\rho_{h_k}(\mathbf{b})$ with $\frac{nr-1}{n-1}$ probability at 1 and exactly $\frac{3nr-2r-n}{n-1}$ at each h_i for $1 \leq i < k$. Thus, the profit from offering prices at most h_{k-1} is $\frac{nr-1}{n-1} - kh(k-1) \frac{3nr-2r-n}{n-1}$. In order to satisfy our lower bound, (4), on $p_{h_k}(\mathbf{b})$, it must put at least $\frac{3nr-2r-n}{n-1}$ on h_k .

Therefore, the probability that the sale price will be no more than kh on masked bid vector on bid vector $\mathbf{b} = (1, \dots, 1, kh, H)$ must be at least $\frac{nr-1}{n-1} + k(\frac{3nr-2r-n}{n-1})$. \square

Given Lemma 1, Theorem 10 is simple to prove.

PROOF. Let \mathcal{A} be a conservative auction. Suppose $\frac{3nr-2r-n}{n-1} = \epsilon > 0$. Let $k = \lceil 1/\epsilon \rceil$, $H \geq kh$, and $h \gg n$. By Lemma 1, $\Pr[f(1, \dots, 1, kh, H) \leq h_k] \geq \frac{nr-1}{n-1} + k\epsilon > 1$. But this is a contradiction, so $\frac{3nr-2r-n}{n-1} \leq 0$. Thus, $r \leq \frac{n}{3n-2}$. The theorem follows. \square

7. CONCLUSIONS AND FUTURE WORK

We have found the optimal auction for the three-unit limited-supply case, and shown that its structure is essentially independent of the benchmark used in its competitive analysis. We have then used this auction to derive the best known auction for the unlimited supply case.

Our work leaves many interesting open questions. We found that the lower bound of [8] is matched by an auction for three bidders,

even when competing against generalized benchmarks. The most interesting open question from our work is whether the lower bound from Theorem 2 can be matched by an auction for more than three bidders. We conjecture that it can.

Second, we consider whether our techniques can be extended to find optimal auctions for greater numbers of bidders. The use of our analytic solution method requires knowledge of a restricted class of auctions which is large enough to contain an optimal auction but small enough that the optimal auction in this class can be found explicitly through analytic methods. No class of auctions which meets these criteria is known even for the four bidder case. Also, when the number of bidders is greater than three, it might be the case that the optimal auction is not expressible in terms of elementary functions.

Another interesting set of open questions concerns aggregation auctions. As we show, the aggregation auction for Υ_3 outperforms the aggregation auction for Υ_2 and it appears that the aggregation auction for Υ_3 is better than Υ_m for $m > 3$. We leave verification of this conjecture for future work. We also show that Υ_3 is not the best three-bidder auction for use in an aggregation auction, but the auction that beats it is able to reduce the competitive ratio of the overall auction only a little bit. It would be interesting to know whether for any m there is an m -aggregation auction that substantially improves on the competitive ratio of Agg_{Υ_m} .

Finally, we remark that very little is known about the structure of the optimal competitive auction. In our auction Υ_3 , the sales price for a given bidder is restricted either to be one of the other bid values or to be higher than all other bid values. The optimal auction for two bidders, the 1-unit Vickrey auction, also falls within this class of auctions, as its sales prices are restricted to bid values. We conjecture that an optimal auction for any number of bidders lies within this class. Our paper provides partial evidence for this conjecture: the lower bound of Section 6 on conservative auctions shows that the optimal auction must offer sales prices higher than any bid value if the lower bound of Theorem 2 is tight, as is conjectured. It remains to show that optimal auctions otherwise only offer sales prices at bid values.

8. ACKNOWLEDGEMENTS

The authors wish to thank Yoav Shoham and Noga Alon for helpful discussions.

9. REFERENCES

- [1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
- [2] S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3, 2003.
- [3] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [4] J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
- [5] A. Fiat, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Competitive generalized auctions. In *Proc. 34th ACM Symposium on the Theory of Computing*, pages 72–81. ACM, 2002.
- [6] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions and digital goods. *Games and Economic Behavior*, 2002. Submitted for publication. An earlier version available as InterTrust Technical Report at URL <http://www.star-lab.com/tr/tr-99-01.html>.

- [7] A. V. Goldberg and J. D. Hartline. Competitiveness via consensus. In *Proc. 14th Symposium on Discrete Algorithms*, pages 215–222. ACM/SIAM, 2003.
- [8] A. V. Goldberg, J. D. Hartline, A. R. Karlin, and M. E. Saks. A lower bound on the competitive ratio of truthful auctions. In *Proc. 21st Symposium on Theoretical Aspects of Computer Science*, pages 644–655. Springer, 2004.
- [9] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proc. 12th Symposium on Discrete Algorithms*, pages 735–744. ACM/SIAM, 2001.
- [10] 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.
- [11] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, 18:511–533, 2001.
- [12] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [13] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proc. of 31st Symp. on Theory of Computing*, pages 129–140. ACM Press, New York, 1999.
- [14] I. Segal. Optimal pricing mechanisms with unknown demand. *American Economic Review*, 16:50929, 2003.
- [15] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.

APPENDIX

A. PROOF OF THEOREM 4

We wish to prove that $\Upsilon_3^{s,t}$, the optimal auction for three bidders against $\mathcal{G}_{s,t}$, has competitive ratio at least $\frac{s^2+t^2}{2t}$. Our proof follows the outline of the proof of Lemma 5 and Theorem 1 from [8]; however, our case is simpler because we only looking for a bound when $n = 3$. Define the random bid vector $\mathbf{B} = (B_1, B_2, B_3)$ with $\Pr[B_i > z] = 1/z$. We compute $\mathbf{E}_{\mathbf{B}}[\mathcal{G}_{s,t}(\mathbf{B})]$ by integrating $\Pr[\mathcal{G}_{s,t}(\mathbf{B}) > z]$. Then we use the fact that no auction can have expected profit greater than 3 on \mathbf{B} to find a lower bound on the competitive ratio against $\mathcal{G}_{s,t}$ for any auction.

For the input distribution \mathbf{B} defined above, let $B_{(i)}$ be the i th largest bid. Define the disjoint events $\mathcal{H}_2 = B_{(2)} \geq z/s \wedge B_{(3)} < z/t$, and $\mathcal{H}_3 = B_{(3)} \geq z/t$. Intuitively, \mathcal{H}_3 corresponds to the event that all three bidders win in $\mathcal{G}_{s,t}$, while \mathcal{H}_2 corresponds to the event that only the top two bidders win. $\mathcal{G}_{s,t}(\mathbf{B})$ will be greater than z if either event occurs:

$$\begin{aligned} \Pr[\mathcal{G}_{s,t}(\mathbf{B}) > z] &= \Pr[\mathcal{H}_2] + \Pr[\mathcal{H}_3] \\ &= 3 \left(\frac{s}{z}\right)^2 \left(1 - \frac{t}{z}\right) + \left(\frac{t}{z}\right)^3 \end{aligned} \quad (5)$$

Using the identity defined for non-negative continuous random variables that $\mathbf{E}[X] = \int_0^\infty \Pr[X > x] dx$, we have

$$\begin{aligned} \mathbf{E}_{\mathbf{B}}[\mathcal{G}_{s,t}(\mathbf{B})] &= t + \int_t^\infty 3 \left(\frac{s}{z}\right)^2 \left(1 - \frac{t}{z}\right) + \left(\frac{t}{z}\right)^3 dz \\ &= 3 \left(\frac{s^2 + t^2}{2t}\right) \end{aligned} \quad (7)$$

Given that, for any auction \mathcal{A} , $\mathbf{E}_{\mathbf{B}}[\mathbf{E}_{\mathcal{A}}[\mathcal{A}(\mathbf{B})]] \leq 3$ [8], it is clear that $\frac{\mathbf{E}_{\mathbf{B}}[\mathcal{G}_{s,t}(\mathbf{B})]}{\mathbf{E}_{\mathbf{B}}[\mathbf{E}_{\mathcal{A}}[\mathcal{A}(\mathbf{B})]]} \geq \frac{s^2+t^2}{2t}$. Therefore, there exists some input \mathbf{b} for each auction \mathcal{A} such that $\frac{\mathcal{G}_{s,t}(\mathbf{b})}{\mathbf{E}_{\mathcal{A}}[\mathcal{A}(\mathbf{b})]} \geq \frac{s^2+t^2}{2t}$.