Competitive Auctions

Andrew V. Goldberg*
goldberg@microsoft.com

Jason D. Hartline†
hartline@cs.washington.edu

Anna R. Karlin†
karlin@cs.washington.edu

Andrew Wright‡
akwright@acm.org

Michael Saks§
saks@math.rutgers.edu

Abstract

We study a class of single-round, sealed-bid auctions for items in unlimited supply, such as digital goods. We introduce the notion of *competitive* auctions. A competitive auction is truthful (i.e., encourages buyers to bid their utility) and yields profit that is roughly within a constant factor of the profit of optimal fixed pricing for all inputs. We justify the use of optimal fixed pricing as a benchmark for evaluating competitive auction profit. We show that several randomized auctions are truthful and competitive and that no truthful deterministic auction is competitive. Our results extend to bounded supply markets, for which we also get truthful and competitive auctions.

A preliminary version of this paper appeared in [6, 7]. The new paper contains new results and refines some of the results and concepts introduced in the preliminary version.

*Microsoft Research, SVC/5, 1065 La Avenida, Mountain View, CA 94043. Part of this work was done while the author was at InterTrust Technologies Corp.

†Computer Science Department, University of Washington. Part of the work done while the author was visiting InterTrust.

‡This work was done while the author was at InterTrust Technologies Corp.

§Department of Mathematics–Hill Center, Rutgers University. 110 Frelinghuysen Rd., Piscataway, NJ 08854.
Supported by NSF grant CCR-9988526. Part of the work was done while visiting Microsoft Research.

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 optimization for the seller in an auction. We focus the design and analysis of truthful, a.k.a., incentive compatible or strategy-proof, auctions that perform well in *worst case* in the presence of unknown market conditions.

The traditional economics approach to the study of profit-maximizing auctions is to construct the optimal Bayesian auction given the prior distribution from which the bidders' utility values, the maximum values that they are willing to pay, are drawn. This has led to a number of interesting results including the use of the Vickrey-Clarke-Groves selling mechanism [4, 8, 14] parameterized by a reservation price that is adjusted to reflect the seller's knowledge of the prior distributions [3, 11]. These results often include very tight bounds on the expected revenue.

In contrast, we focus on the design of unparameterized auctions for profit maximization in unknown markets, i.e., when Bayesian priors are not known.[1] There are a number of compelling reasons for considering profit maximization from this perspective. In order to use a Bayesian optimal auction the prior distribution, or at least a reasonable approximation thereof, must be ascertained in advance. In some cases, determining the prior distribution in advance may not be convenient or possible. Moreover, even if prior distributions are known, when many auctions are to be run, computing and setting a new reservation price for each one may be inconvenient or infeasible.

The study of profit maximizing auctions for unknown markets provides understanding for auction problems where the seller has incomplete knowledge of the distribution of bidder valuations. We can determine how well an uninformed mechanism can perform. We can get a quantitative understanding of the relative value of being informed. A partially informed auctioneer choosing an auction mechanism has to consider the trade-off between using an auction tailored to assumptions about bidder valuations that may or may not be correct versus using an auction designed to work as well as possible under unknown market conditions.

Worst case analysis is the only way to give provable performance guarantees in the presence of total uncertainty. As such, we will use worst case analysis to measure the performance of uninformed auctions. We will show that it is possible to obtain optimal performance, to within a multiplicative constant factor, even under arbitrary, unknown and worst-case market conditions.

In the case where the auctioneer has partial information on possible inputs, it is possible to tailor the analysis and prove better bounds than are possible in the worst case over all possible inputs. As an example, we consider the performance of our auctions in the *mass market* case, where the auctioneer expects to sell many items. In this case, we give an analysis of our auctions for the worst input over a practical restriction of bidders' valuations that is relevant to mass market

[1] In the literature, such auctions are sometimes called *detail-free* auctions.

sales. Similarly, though we do not provide the analysis in this paper, a Bayesian analysis of our auctions could be performed to obtain tighter bounds for the Bayesian setting. An ideal auction has good performance in worst case and a better performance when the inputs come from specific prior distributions.

To analyze auction performance without making assumptions about the prior input distribution, a criterion is needed for gauging auction revenues. Our approach is motivated by *competitive analysis* of on-line algorithms [2, 13], where the performance of an on-line algorithm, one that must make decisions without knowledge of the future, is compared to the performance of an optimal off-line algorithm that knows the future. The assumption that the on-line algorithm not know the future in advance is analogous to the assumption that auctions not know the bidders' values in advance. We gauge a truthful auction's performance on a particular bid set by comparing it against the profit that would be achieved by an "optimal" omniscient auction, one that knows the bids in advance, on the same set of bids. An auction is *competitive* if it achieves a profit that is a constant fraction of optimal *on every input*.

In this paper, we focus on the case where the auctioneer has an unlimited number of indivisible items available to sell, each consumer wants at most one item, and the auctioneer has no value for the items. We use the term *unlimited supply* to describe this supply curve. This unlimited supply case 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. All of our results generalize naturally to the *bounded supply* case where the auctioneer has a limited number of items to sell (fewer than the number of bidders). Adhering to the *revelation principle* [11], we will restrict our attention to single-round, sealed-bid, truthful auction mechanisms. In truthful auctions, truth-telling, i.e, revealing their true utility value as their bid, is a dominant strategy for each bidder.

The rest of the paper is organized as follows. In Section 2 we formally define auction mechanisms and the bidder model. We also review an algorithmic characterization of deterministic and randomized truthful mechanisms.

In Section 3 we describe the basics of competitive analysis and motivate the its use for auction problems by contrasting it to Bayesian analysis. We also discuss in detail the competitive framework that we will be using in this paper for both the worst case and mass market case.

In Section 4 we show that no deterministic auction has good worst case performance. This motivates the use of randomized auction mechanisms. In Section 5 we give a bound on how well any auction can perform in worst case. We give two randomized *competitive* auctions in Section 6 and prove that they perform well in worst case. The first auction uses the optimal price for a random sample of the bids as the sale price for the remaining bids, i.e., does market analysis on the sample. Indeed, we will show that for mass markets, this auction obtains a near optimal profit. Moreover, it is performs within a constant factor of optimal in our worst case competitive framework. The second auction we introduce is based on a standard cost sharing mechanism. Analysis of this auction is much simpler than the first and its worst case performance is much better.

In Section 7 we discuss the extension of the results for unlimited supply to the limited supply case where there are fewer items for sale than bidders. In Section 8 we discuss a natural subclass of auctions for profit maximization which includes all of our auctions. We conclude in Section 9.

2 Preliminaries and Notation

We consider single-round, sealed-bid auctions for a set of identical items available in unlimited supply.

Definition 2.1 *A* single-round sealed-bid auction*, $\mathcal{A}$, is one where:*

1. *Each bidder submits a bid, representing the maximum amount they are willing to pay for an item. We denote by* $\mathbf{b}$ *the vector of all submitted bids, i.e., the* input*. The i-th component of* $\mathbf{b}$ *is b_i, the bid submitted by bidder i. We denote by n the number of bidders.*

2. *Given the bid vector* $\mathbf{b}$*, the auctioneer computes an* output *consisting of an* allocation*,* $\mathbf{x} = (x_1, \ldots, x_n)$*, and* prices*,* $\mathbf{p} = (p_1, \ldots, p_n)$*. The allocation x_i is an indicator for bidder i's receipt of the item (1 if bidder i receives the item and 0 otherwise). If $x_i = 1$, we say that* bidder i *wins*. Otherwise, bidder i *loses*, or is rejected*. The price, p_i, is what bidder i pays the auctioneer. We assume that $0 \leq p_i \leq b_i$ for all winning bidders and that $p_i = 0$ for all losing bidders (these are the standard assumptions of* no positive transfers *and* voluntary participation*. See, e.g., [10]).*

3. *The* profit *of the auction (or auctioneer) is* $\mathcal{A}(\mathbf{b}) = \sum_i p_i$*.*

We say that an auction is *deterministic* if the allocation and prices are completely determined as a function of the bid vector. We say that the auction is *randomized* if the procedure by which the auctioneer computes the allocation and prices is randomized. Note that if the auction is randomized, the profit of the auction, the output prices, and the allocation are random variables.

We make the following assumptions about bidders:

- Each bidder has a private utility value, representing the true maximum they are willing to pay for an item. We denote by u_i bidder i's utility value.
- Each bidder bids so as to maximize their *profit*, $u_i x_i - p_i$.
- Bidders bid with full knowledge of the auctioneer's mechanism.
- Bidders do not collude.
- The bidders in the auction are indistinguishable from the perspective of the auctioneer. Thus, the auctioneer cannot perform any knowledgeable market segmentation.

Finally, we formally define the notion of truthfulness.

Definition 2.2 *We say that a deterministic auction is* truthful *if, for each bidder i and any choice of bid values for all other bidders, bidder i's profit is maximized by bidding their utility value, i.e., by setting $b_i = u_i$.*

Thus, we consider an *ex post* notion of truthfulness with respect to bidders strategies. Truthful auctions encourage rational bidders to bid their utility value if the bid value that maximizes their profit is unique. If the bid value is not unique, truthfulness at least does not discourage bidders from bidding their utility value.

For randomized auctions, one natural approach is to say that an auction is truthful if bidding utility maximizes a bidder's expected profit. We use another natural, but stronger notion of truthfulness that allows us to take advantage of a particularly useful characterization of truthful auctions, discussed in Section 2.1. We require randomized auctions to be truthful *ex post* with respect to the randomness in the mechanism.

Definition 2.3 *A randomized auction is* truthful *if it can be described as a probability distribution over deterministic truthful auctions.*

Clearly, with this notion of truthfulness, the probability that a bidder's profit exceeds v is simultaneously maximized for every v by bidding truthfully.

Henceforth we will consider only truthful auctions. As bidding u_i is a dominant strategy for bidder i in a truthful auction, in the remainder of this paper, we assume that $b_i = u_i$.

2.1 Bid-Independent Auctions

Intuitively, if a price a bidder gets in an auction is independent of the bidder's bid value, the auction is truthful. Bid-independent auctions formalize this observation and give a characterization of truthful auctions. Various communities have formulations of this characterization. However, as we rely heavily upon this result, we give an elementary proof.

We begin by formalizing the notion of bid-independence. Let $\mathbf{b}_{-i}$ denote the vector of bids $\mathbf{b}$ with b_i removed, i.e., $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n)$. We call such a vector *masked.* Let f be a function from masked vectors to prices (non-negative real numbers) The *deterministic bid-independent auction defined by* f is $\mathcal{A}_f$ (Auction 1).

Auction 1 Bid-independent Auction: $\mathcal{A}_f(\mathbf{b})$

For each bidder i:

1. $t_i \leftarrow f(\mathbf{b}_{-i})$.
2. If $t_i \leq b_i$, set $x_i \leftarrow 1$ and $p_i \leftarrow t_i$. (Bidder i wins.)
3. Otherwise, set $x_i = p_i = 0$. (In this case, we say that bidder i is rejected.)

In step 2 above, as the proof of Theorem 2.4 shows, bid-independence allows the inequality, $t_i \leq b_i$, to be strict or non-strict at the discretion of $f(\mathbf{b}_{-i})$. Thus f can specify whether to accept b_i at price t_i if b_i is in (t_i, ∞) or $[t_i, \infty)$.

For example, by setting $f = \max$ for all i and breaking ties arbitrarily, we obtain the 1-item Vickrey auction, in which the highest bidder wins at the second highest price. Similarly, if f is the function that returns the k-th highest bid, we get the k-item Vickrey auction.

Theorem 2.4 *A deterministic auction is truthful if and only if it is equivalent to a deterministic bid-independent auction.*

The theorem follows from the following two lemmas.

Lemma 2.5 *Any deterministic bid-independent auction is truthful.*

Proof. Consider the auction outcome for bidder i. If they bid at least t_i they win and pay t_i; otherwise they lose. If their utility is below t_i, there is no way they can win and have a positive profit. If $u_i \geq t_i$, then any bid of at least t_i allows them to win the auction and pay t_i; bidding below t_i would cause them to lose. Thus, for any utility value and any t_i, bidding u_i maximizes the bidder i's profit. ■

The following result completes the proof of equivalence of truthfulness and bid-independence for deterministic auctions.

Lemma 2.6 *Any truthful deterministic auction is equivalent to a deterministic bid-independent auction.*

Proof. Given any truthful deterministic $\mathcal{A}$ we can determine an f such that the bid-independent implementation, $\mathcal{A}_f$, is identical to $\mathcal{A}$. Let $\mathbf{b}_i^x = (b_1, \ldots, b_{i-1}, x, b_{i+1}, \ldots, b_n)$, the bid vector obtained by replacing b_i with x. If there is some value x^* such that in $\mathcal{A}(\mathbf{b}_i^{x^*})$ bidder i wins and pays p (note this requires $p \leq x^*$) then define $f(\mathbf{b}_{-i})$ to be p. $f(\mathbf{b}_{-i})$ uses a closed interval $[p, \infty)$ for winners if $\mathcal{A}(\mathbf{b}_i^p) = p$ and an open interval otherwise.

Give this value of p, we now show for $\mathcal{A}(\mathbf{b}_i^x)$ that:

1. If bidder i wins, he pays p.

2. Bidder i wins by bidding any value $x > p$ (and possibly by bidding $x = p$).

To see 1, assume to the contrary that there is some other bid value y such that running $\mathcal{A}(\mathbf{b}_i^y)$ results in bidder i winning and paying $q \neq p$. Without loss of generality $q > p$ so a bidder with utility y would have a higher profit by bidding x^*. This contradicts $\mathcal{A}$'s truthfulness.

To see 2, assume to the contrary that there is some bid value $y \in (p, \infty)$ such that bidder i does not win by bidding y. Notice that a bidder with utility y would have a higher profit by bidding x^*. Again this contradicts the $\mathcal{A}$'s truthfulness and gives the lemma. ■

Definition 2.7 *A* randomized bid-independent auction *is a probability distribution over bid-independent auctions. For these auctions, $f(\mathbf{b}_{-i})$ is a non-negative real-valued random variable.*

Note that the random variables $f(\mathbf{b}_{-i})$ and $f(\mathbf{b}_{-j})$ need not be independent. It follows immediately from this definition, Definition 2.3, and Theorem 2.4 that:

Corollary 2.8 *A randomized auction is truthful if and only if it is equivalent to a randomized bid-independent auction.*

2.2 Optimal Omniscient Auctions

In our discussions to follow, it will be useful to compare the performance of truthful auctions to that of the optimal single price and the optimal multiple-price omniscient auctions.

Definition 2.9 *The* optimal single price omniscient auction*, $\mathcal{F}$, is defined as follows: Let* $\mathbf{b}$ *be a bid vector, and let v_i be the i-th largest bid in* $\mathbf{b}$*. Auction $\mathcal{F}$ on input* $\mathbf{b}$ *determines the value k such that kv_k is maximized. All bidders with $b_i \geq v_k$ win at price v_k; all remaining bidders lose. The profit of $\mathcal{F}$ on input* $\mathbf{b}$ *is thus*

$$\mathcal{F}(\mathbf{b}) = \max_{1 \leq i \leq n} iv_i.$$

Definition 2.10 *The* optimal multiple price omniscient auction*, $\mathcal{T}$, is the auction that sells to each bidder at his bid value. Thus, the profit of $\mathcal{T}$ on input* $\mathbf{b}$ *is*

$$\mathcal{T}(\mathbf{b}) = \sum_{1 \leq i \leq n} b_i.$$

3 Competitive Analysis

The goal of this section is to motivate and explain the application of worst-case competitive analysis to auctions. We begin with an example of an auction that performs well under a Bayesian analysis. Then we show that under worst-case analysis, the same auction does poorly. Finally, we describe how competitive analysis can be used to give stronger results.

We begin by considering a natural truthful auction, the *deterministic optimal threshold* (DOT) auction. To define DOT, we introduce the notion of the optimal sale price for a set of bids.

Definition 3.1 *Let* $\mathbf{b}$ *be a vector of bids. Denote by* $\text{opt}(\mathbf{b})$ *the sale price for* $\mathbf{b}$ *that gives the optimal profit, i.e.,*

$$\text{opt}(\mathbf{b}) = \text{argmax}_{v_i}\, iv_i,$$

where v_i is the i-th largest bid in $\mathbf{b}$*.*

DOT is the truthful auction defined as follows:

Definition 3.2 *The* deterministic optimal threshold (DOT) *auction is defined by the bid independent function f :*

$$f(\mathbf{b}_{-i}) = \text{opt}(\mathbf{b}_{-i}).$$

In other words, DOT uses for each bid b_i the optimal sale price for the remaining bids $\mathbf{b}_{-i}$*.*

Note that the optimal price for $\mathbf{b}$ could be the same as the optimal price for $\mathbf{b}_{-i}$ but in general this is not the case. The following lemma is not difficult to prove.

Lemma 3.3 *Assume that the bids in* $\mathbf{b}$ *are i.i.d. from any bounded support probability distribution. Then, as the number of bidders tends to infinity, the profit of the DOT auction converges to the optimal profit.*

Thus, DOT is a profit maximizing auction for a large class of distributions over bid vectors: any bounded support i.i.d. distribution. However, it is also easy to exhibit classes of bid vectors where DOT's profit is very far from optimal.

For example, consider n bidders where n/b of them bid $b \gg 1$ and the remaining bidders bid 1. Running DOT on this bid vector will result in the following: For a bid at value b, we remove it and compute opt of the remaining bids. Of the $n-1$ bids remaining $n/b-1$ of them are at b. Thus, opt outputs 1 as $n-1$ bids at price 1 gives a higher revenue than $n/b-1$ bids at price b. Similarly for a bid at value 1, we remove it and compute opt of the remaining bids. Of the $n-1$ of them, there are n/b of them at b, the rest at 1. Thus, opt chooses to output b because n/b bids at price b gives a higher revenue than $n-1$ bids at price 1. Thus, all bids at value 1 are rejected and all bids at value b win the auction and only have to pay 1. DOT's profit is thus n/b (the number of bids at b) whereas the optimal single price profit is n. Thus, for b very large, DOT performs arbitrarily worse than the optimal single price profit.

We note that any input on which the single price omniscient auction finds two different bid values that give approximately the same revenue will have a similar worst case. Indeed, we extend this result in Section 4 by showing that all deterministic auctions suffer from this type of problem.

In the Bayesian analysis of Lemma 3.3, these bad inputs are very unlikely to occur; analysis of the expected profit hides the possibility of messing up on specific inputs. While this is not a problem if the bids are from a random source, it is not necessarily reasonable to expect the bidders' utility values to be truly random.

To expose auctions, like the one above, that have inputs on which they perform poorly, we turn to worst-case analysis. The goal of the worst-case analysis we employ here is to get a bound on how badly the auction could mess up if an adversary generated the worst possible input. While this adversarial view is clearly not a realistic model of how bidders bid either, an auction that performs well in the worst case will certainly perform well on the more benign reality. As we shall see in this paper, *there are auctions that perform well even under worst-case analysis.*

The notion of competitive analysis is a natural one. In the worst case, all input bids could be zero (or negligible) and no auction will be able to extract a high profit. Thus, instead of absolute profit, we look at auction profit relative to an "optimal" auction on the same input. The correct notion of an optimal auction is not obvious. In general, our goal is to get the strongest results we can, and thus we will try to compare our truthful auctions against the best possible auction that we can feasibly compete with. Once the optimal auction is defined, the best *competitive auction* is the one that minimizes the *competitive ratio*, the worst case ratio between the optimal auction profit and its profit. A fundamental goal of this research is to determine the best metric for comparison, the corresponding best competitive ratio, and the auction mechanism that achieves it.

3.1 Competitive Auction Framework

A key part of setting up a competitive framework for analyzing solutions to any problem is coming up with the right metric for comparison. As a starting point, we would like to take the strongest possible benchmark for comparison that we can: the profit of an auctioneer that is perfectly informed about the bidder's values. This leads us to consider as the two most natural metrics for comparison the optimal omniscient multi-price and single-price auctions, $\mathcal{T}$ and $\mathcal{F}$, defined in Section 2.2.

We first compare the performance of $\mathcal{F}$ relative to $\mathcal{T}$. Specifically, we observe that in the worst case, the maximum ratio of $\mathcal{T}$ to $\mathcal{F}$ is logarithmic in the number n of bidders.

Lemma 3.4 *There exist bid vectors* $\mathbf{b}$ *for which*

$$\mathcal{F}(\mathbf{b}) = \Theta(\mathcal{T}(\mathbf{b})/\ln n).$$

Moreover, for all bid vectors $\mathbf{b}$

$$\mathcal{F}(\mathbf{b}) \geq \mathcal{T}(\mathbf{b})/\ln n.$$

Proof. For the first part, let $\mathbf{b}$ be the bid vector such that $b_i = n/i$. Then $\mathcal{F}(\mathbf{b}) = n$ and $\mathcal{T}(\mathbf{b}) = n(\ln(n) + \Theta(1))$.

For the second part, let v_i be the i-th largest bid in bid vector $\mathbf{b}$. Suppose that $\mathcal{F}(\mathbf{b}) = \max_i iv_i = kv_k$. Then for all i,

$$iv_i \leq kv_k.$$

Thus,

$$\mathcal{T}(\mathbf{b}) = \sum_{i=1}^{n} v_i \leq \sum_{i=1}^{n} \frac{kv_k}{i} \leq \mathcal{F}(\mathbf{b}) \sum_{j=1}^{n} \frac{1}{j} = \mathcal{F}(\mathbf{b})(\ln n + O(1)).$$

■

Now we show that no truthful auction can be competitive against $\mathcal{F}$ (and hence it can not be competitive against $\mathcal{T}$).

Lemma 3.5 *For any truthful auction* $\mathcal{A}_f$ *and any* $\beta \geq 1$*, there is a bid vector* $\mathbf{b}$ *such that the expected profit of* $\mathcal{A}_f$ *on* $\mathbf{b}$ *is less than* $\mathcal{F}(\mathbf{b})/\beta$*.*

Proof. Consider a bid-independent randomized auction on two bids, 1 and $x \geq 1$. Let h be the smallest value greater or equal to 1 such that $\mathbf{Pr}[f(1) \geq h] \leq \frac{1}{2\beta}$. Then the profit on input vector $\mathbf{b} = (1, H)$ with $H = 4\beta h$ is at most

$$\frac{H}{2\beta} + h(1 - \frac{1}{2\beta}) + 1 < 4h = \frac{H}{\beta} = \frac{\mathcal{F}(\mathbf{b})}{\beta}.$$

■

Lemma 3.5 shows that we cannot expect to come close to matching the performance of the optimal single price omniscient auction in the case where the optimal profit is generated from the single highest bid. Thus, we must set our sights slightly lower. Later in the paper we will present auctions that are *competitive* with $\mathcal{F}^{(2)}$, the optimal single price auction that sells at least two items. Such auctions perform comparably to $\mathcal{F}^{(2)}$ in that they achieve a constant fraction of the revenue of $\mathcal{F}^{(2)}$ on *all inputs.*

Definition 3.6 *The* optimal single price omniscient auction that sells at least two items, $\mathcal{F}^{(2)}$*, is defined as follows: Let* $\mathbf{b}$ *be a bid vector, and let* v_i *be the* i*-th largest bid in the vector* $\mathbf{b}$*. Auction* $\mathcal{F}^{(2)}$ *on input* $\mathbf{b}$ *determines the value* k *such that* $k \geq 2$ *and* kv_k *is maximized. All bidders with* $b_i \geq v_k$ *win at price* v_k*; all remaining bidders lose. The profit of* $\mathcal{F}^{(2)}$ *on input* $\mathbf{b}$ *is thus*

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

Note that for $\mathbf{b}$ where $\mathcal{F}$ elects to sell at least two items, $\mathcal{F}^{(2)}(\mathbf{b}) = \mathcal{F}(\mathbf{b})$. Thus, excluding bid vectors where only the highest bidder wins in the optimal auction, comparing auction performance to $\mathcal{F}^{(2)}$ is identical to comparing it to $\mathcal{F}$.

Next we generalize the definition of $\mathcal{F}^{(2)}$ to define $\mathcal{F}^{(m)}$, as the optimal single price omniscient auction that sells at least m items, and formalize our definition of competitiveness.

Definition 3.7 *The* m-optimal single price omniscient auction $\mathcal{F}^{(m)}$ *is defined as follows: Let* $\mathbf{b}$ *be a bid vector, and let* v_i *be the* i*-th largest bid in the vector* $\mathbf{b}$*. Auction* $\mathcal{F}^{(m)}$ *on input* $\mathbf{b}$ *determines the value* k *such that* $k \geq m$ *and* kv_k *is maximized. All bidders with* $b_i \geq v_k$ *win at price* v_k*; all remaining bidders lose. The profit of* $\mathcal{F}^{(m)}$ *on input* $\mathbf{b}$ *is thus*

$$\mathcal{F}^{(m)}(\mathbf{b}) = \max_{m \leq k \leq n} kv_k.$$

Finally, we formalize the notion of a competitive auction.

Definition 3.8 *We say that auction* $\mathcal{A}$ *is* β-competitive against $\mathcal{F}^{(m)}$ *if for all bid vectors* $\mathbf{b}$*, the expected profit of* $\mathcal{A}$ *on* $\mathbf{b}$ *satisfies*

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

We say that an auction is competitive against $\mathcal{F}^{(m)}$ *if the auction is* β*-competitive against* $\mathcal{F}^{(m)}$ *for a constant* β*. We refer to* β *as the* competitive ratio *of* $\mathcal{A}$*.*

The statement that an auction is competitive against $\mathcal{F}^{(m)}$ implies that, restricted to inputs $\mathbf{b}$ such that there are at least m items sold by the optimal auction, i.e., $\mathcal{F}^{(m)}(\mathbf{b}) = \mathcal{F}(\mathbf{b})$, our auctions are competitive against $\mathcal{F}$. For $m = 2$, this restriction precisely excludes the case where it is not possible to be competitive – when there is one bidder with very large utility.

Observe that since the profit of $\mathcal{F}^{(m)}$ decreases as m increases, as we compete against $\mathcal{F}^{(m)}$ for larger values of m, we are demanding less and less of the truthful auction. We will thus single out the case of competition against $m = 2$, as $\mathcal{F}^{(2)}$ is the strongest omniscient auction that we will be able to feasibly compete with.

Definition 3.9 *We say an auction is* β-competitive *if it is* β*-competitive against* $\mathcal{F}^{(2)}$*. In cases where we do not wish to specify the constant* β*, we simply say that the auction is* competitive.

By considering competitiveness against $\mathcal{F}^{(m)}$ for different values of m, we obtain results that are relevant to a wider range of applications. For example, in situations in which the auctioneer does not have prior distributions on bidders' bid values, but nonetheless has a limited amount of prior knowledge about the bidders, i.e., that his profit will be maximized by selling at least m items, he can use an auction that is tailored to such a situation, and obtain much stronger guarantees on the competitive ratio. One such natural example is when there are a large number of bidders, and the auctioneer is safe in assuming that all bid values come from a bounded range.

4 Deterministic Auctions are not Competitive

In this section, we show that no symmetric deterministic auction can be competitive. An auction is *symmetric* if the outcome is independent of the order of the bids. More precisely, we say that $\mathcal{A}$ is symmetric if for all bid vectors $\mathbf{b}$ and permutations π of the bidders, the output of $\mathcal{A}$ on input $\pi(\mathbf{b})$ is price vector $\pi(\mathbf{p})$ and allocation $\pi(\mathbf{x})$ (given that the output of $\mathcal{A}$ on $\mathbf{b}$ is $\mathbf{p}$ and $\mathbf{x}$).

We now show that no symmetric deterministic auction is competitive. In contrast, in Section 6, we will show that there are competitive symmetric randomized auctions.

Theorem 4.1 *Let $\mathcal{A}_f$ be any symmetric deterministic auction defined by bid-independent function f. Then $\mathcal{A}_f$ is not competitive: For any $1 \geq m \leq n$ there exists a bid vector $\mathbf{b}$ of length n such that the profit of $\mathcal{A}_f$ on $\mathbf{b}$ is at most $\mathcal{F}^{(m)}(\mathbf{b})\frac{m}{n}$.*

Proof. Fix n and m and the symmetric bid-independent auction $\mathcal{A}_f$. Consider the set of bid vectors whose bids are all n or 1. For $0 \leq j \leq n-1$, write $f(j)$ for the price the auction assigns to a masked vector with exactly j bids at n and $n-1-j$ bids at 1. Note that if $f(0) > 1$ then if all bids are 1, the auction has profit 0 and the conclusion of the theorem holds trivially. So assume $f(0) \leq 1$ and let k be the largest integer in $\{0, \ldots, n-1\}$ such that $f(k) \leq 1$. Let $\mathbf{b}$ be the the bid vector with $k+1$ bids at n and $n-k-1$ bids at 1. The profit of $\mathcal{A}_f$ on $\mathbf{b}$ is $(k+1)f(k) \leq k+1$ since the only winners are those who bid n. If $k \leq m-1$ then $\mathcal{F}^{(m)}$ has profit n. If $k \geq m$ then $\mathcal{F}^{(m)}$ has profit at least $(k+1)n$. In either case the profit is at most $\mathcal{F}^{(m)}(\mathbf{b})\frac{m}{n}$. ■

This result motivates the consideration of randomized mechanisms. As a randomized mechanism is just a randomization over deterministic auctions, it is still the case that a randomized auction can perform poorly on unlucky outcomes of our random coin flips. However, the auction can be designed such that these unlucky outcomes are very improbable. We note that this use of randomness in the mechanism is very different from assuming the bids come from a random distribution. The latter is making an external assumption on our inputs, while internal randomness in the mechanism is guaranteed and completely under the control of the mechanism. As we will show with our development of randomized auctions that are competitive in worst case, there is no need to make the assumption that the bids are randomly generated when we can randomize the mechanism instead. Of course, the use of randomization in mechanisms and algorithms is a standard technique in game theory and computer science

5 A Lower Bound on the Competitive Ratio

We have seen that in order to construct competitive truthful auctions, we will need to incorporate randomization into the mechanism. Before showing how this is done, we prove a lower bound on the performance of any truthful auction, even a randomized one, compared to $\mathcal{F}^{(2)}$: We show that for any randomized truthful auction $\mathcal{A}$, there exists an input bid vector $\mathbf{b}$ on which

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

.

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 show that $\mathbf{E}_{\mathbf{b}}[\mathbf{E}_{\mathcal{A}}[\mathcal{A}(\mathbf{b})]] \leq \frac{\mathbf{E}_{\mathbf{b}}\left[\mathcal{F}^{(2)}(\mathbf{b})\right]}{2.42}$. It follows immediately (from the definition of $\mathbf{E}_{\mathbf{b}}[\cdot]$) that there must exist a fixed choice of $\mathbf{b}$ (depending on $\mathcal{A}$) for which $\mathbf{E}[\mathcal{A}(\mathbf{b})] \leq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{2.42}$.

Consider n i.i.d. bids $\mathbf{b}^{(n)}$ generated from the distribution with each bid b_i satisfying $\mathbf{Pr}[b_i > y] = 1/y$ for all $y \geq 1$. Consider a truthful auction $\mathcal{A}$. Let V_i be the price offered to b_i in the bid-independent implementation of $\mathcal{A}$. V_i is a random variable depending on $\mathcal{A}$ and all of the b_j other than b_i. Let P_i be the profit from bidder i which is 0 if $b_i < V_i$ and V_i otherwise. For $v \geq 0$, $\mathbf{E}[P_i | V_i = v] = v \cdot \mathbf{Pr}[b_i > v | V_i = v] = v \cdot \mathbf{Pr}[b_i > v] \leq 1$, since b_i is independent of V_i. Therefore $\mathbf{E}[P_i] \leq 1$. Thus we have:

Lemma 5.1 *For $\mathbf{b}^{(n)}$ defined above, the expected revenue of any truthful deterministic auction, $\mathcal{A}$, is at most n.*

Note that for any deterministic bid-independent auction that offers prices of at least one, the expected revenue is exactly n.

The proof of the following result is technical and can be found in Appendix C.

Lemma 5.2 *For n bids from the above distribution, the expected value of $\mathcal{F}^{(2)}$ is*

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

Combining this with the previous lemma we get:

Lemma 5.3 *For bids as defined above, we have*

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

Interesting special cases are $n = 2$ where this gives a lower bound of 2 which matches the best competitive auction for two bids, the 1-item Vickrey auction. For $n = 3$ this gives a lower bound of 13/6. A lower bound for the ratio of the best competitive auction on general n is obtained by taking the limit.

Lemma 5.4

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

The proof of the lemma appears in Appendix C. Routine calculations then show that the limit value is at least 2.42.

Combining this with Yao's application of the Minimax theorem (See Motwani and Raghavan [9], Section 2.2.), we get:

Theorem 5.5 *Let $\mathcal{A}$ be any truthful randomized auction. The competitive ratio of $\mathcal{A}$ is at least* 2.42.

6 Competitive Auctions via Random Sampling

We now present two techniques for designing competitive auctions that are based on random sampling. Both techniques start by randomly partitioning the input bids into two sets. This is done by flipping a fair coin for each bid to decide which partition to assign it. We then use one partition for market analysis to plug into a sub-auction to be run on the other partition, and vice versa. In what follows, we explore two choices for the sub-auction. The first, the Dual-Price Sampling Optimal Threshold auction, uses as a sub-auction a simplification of the optimal Bayesian auction (Vickrey with reservation price) to the unlimited supply case. The second, the Sampling Cost Sharing auction, uses the cost sharing mechanism of Moulin and Shenker [10] for the sub-auction.

6.1 Dual-Price Sampling Optimal Threshold Auction

In this section, we present the Dual-Price Sampling Optimal Threshold (DSOT) auction. DSOT is a randomized version of the DOT auction presented in Section 3. The DSOT auction is guaranteed, on all inputs, to achieve a constant fraction of the profit of $\mathcal{F}^{(2)}$, the optimal single-price omniscient auction which sells at least two items. More importantly, on a large class of interesting and practical bid vectors, DSOT is guaranteed to get very close the profit of $\mathcal{F}$, the optimal single-price omniscient auction.

The DSOT auction is defined as follows:

Note that DSOT is a dual-priced auction. If it is important to have a single-price auction, one can simply skip the 4-th step and reject all the bids in $\mathbf{b}'$, at the cost of half the expected profit.

It is readily apparent that DSOT can be implemented bid-independently and thus from Corollary 2.8:

Observation 6.1 *The Dual-Price Sampling Optimal Threshold auction is truthful.*

We now discuss the performance of DSOT. Our first result is that on every input, DSOT achieves a constant fraction of the profit of $\mathcal{F}^{(2)}$.

Theorem 6.2 DSOT *is constant competitive against $\mathcal{F}^{(2)}$.*

Auction 2 Dual-Price Sampling Optimal Threshold Auction (DSOT)

1. Partition bids $\mathbf{b}$ uniformly at random into two sets: for each bid, with probability 1/2 put the bid in $\mathbf{b}'$ and otherwise $\mathbf{b}''$.
2. Let $p' = \mathrm{opt}(\mathbf{b}')$ and $p'' = \mathrm{opt}(\mathbf{b}'')$, the optimal fixed price thresholds for $\mathbf{b}'$ and $\mathbf{b}''$, respectively. (See Definition 3.1.)
3. Use p' as a threshold for all bids in $\mathbf{b}''$ (i.e., all bids in $\mathbf{b}''$ of value below p' are rejected; all remaining bids win at price p').
4. Use p'' as a threshold for all bids in $\mathbf{b}'$.

Proofs of this theorem and the rest of the theorems in this section are technical and thus are deferred to Appendix A.

The constant bound we obtain in this theorem is quite weak. However, there are a number of interesting special cases in which DSOT's performance is significantly better. One such special case is presented in the following theorem.

Theorem 6.3 *Let* $\mathbf{b}$ *be any bounded-range bid vector, i.e., any bid vector of* n *bids with* $b_i \in [1, h]$ *for all* i*. Then*

$$\lim_{n\to\infty} \max_{\mathbf{b}} \frac{\mathcal{F}(\mathbf{b})}{\mathrm{DSOT}(\mathbf{b})} = 1.$$

To prove this theorem, and generalizations thereof, we consider the DSOT_r auction, a parameterized version of the DSOT auction. To define DSOT_r, we first generalize Definition 3.1.

Definition 6.4 *Let* b *be a vector of bids. Denote by* $\mathrm{opt}_r(\mathbf{b})$ *the sale price for* $\mathbf{b}$ *that gives the optimal profit among those sale prices that result in the sale of at least* r *items, i.e.,*

$$\mathrm{opt}_r(\mathbf{b}) = \mathrm{argmax}_{v_i | i \geq r}\, i v_i,$$

where v_i *is the* i*-th largest bid in* $\mathbf{b}$*. If* $r > n$*, we arbitrarily define* $\mathrm{opt}_r(\mathbf{b}) = 0$*.*

Theorem 6.5 *There is an absolute constant* C*, such that for any* $\epsilon > 0$*,* $\mathrm{DSOT}_{\frac{m}{2}-\epsilon m}$ *is* $(1+\epsilon)$*-competitive against* $\mathcal{F}^{(m)}$*, with probability at least* $1 - e^{-C\epsilon^2 m}$*.*

Thus, a parameterized version of DSOT asymptotically (as m gets large) matches the profit of $\mathcal{F}^{(m)}$. Recall that $\mathcal{F}^{(m)}$, the optimal single-price auction that is required to have at least m winners, is the same as $\mathcal{F}$ when $\mathcal{F}$ chooses to sell at least m items. Thus $\mathcal{F}^{(m)}$ is the best auction when the auctioneer is required to use a single price and sell at least m items. An immediate corollary of the above theorem is that on any bid vector $\mathbf{b}$ such that (a) $\mathcal{F}(\mathbf{b}) = \mathcal{F}^{(m)}(\mathbf{b})$ and (b) with high probability DSOT and $\mathrm{DSOT}_{\frac{m}{2}-\epsilon m}$ exhibit the same behavior on $\mathbf{b}$ (i.e., for $\mathbf{b}$ partitioned into $\mathbf{b}'$

Auction 3 Parameterized Dual-Price Sampling Optimal Threshold Auction (DSOT_r)

1. Partition bids $\mathbf{b}$ uniformly at random into two sets: for each bid, with probability 1/2 put the bid in $\mathbf{b}'$ and otherwise $\mathbf{b}''$.
2. Let $p' = \text{opt}_r(\mathbf{b}')$ and $p'' = \text{opt}_r(\mathbf{b}'')$, the optimal fixed price thresholds that sell at least r items for $\mathbf{b}'$ and $\mathbf{b}''$, respectively.
3. Use p' as a threshold for all bids in $\mathbf{b}''$ (i.e., all bids in $\mathbf{b}''$ of value below p' are rejected; all remaining bids win at price p').
4. Use p'' as a threshold for all bids in $\mathbf{b}'$.

---

and $\mathbf{b}''$, $\text{opt}_r(\mathbf{b}') = \text{opt}(\mathbf{b}')$ and $\text{opt}_r(\mathbf{b}'') = \text{opt}(\mathbf{b}'')$, where $r = \frac{m}{2} - \epsilon m)$, we can conclude that with high probability $\text{DSOT}(\mathbf{b}) \geq \frac{\mathcal{F}(\mathbf{b})}{(1+\epsilon)}$.

It is easy to check that conditions (a) and (b) hold for the case of bids of bounded support, as the number of bidders gets large, yielding Theorem 6.3 as an immediate corollary of Theorem 6.5.

Thus, for a large class of inputs, DSOT achieves essentially optimal performance. For worst-case inputs, however, the constant in the competitiveness of DSOT is weak. It is quite clear from studying the proof of Theorem 6.2 (in Appendix A) that the analysis given there is not tight and the constant bound we obtain on DSOT's competitiveness is very weak. However, it is not hard to see that DSOT can not be better than 4-competitive. For example, when $\mathbf{b}$ consist of two very high bids h and $h + \epsilon$, and all other bids are negligible, the expected profit of DSOT $= \mathcal{F}^{(2)}/4$. In the next section, we present an auction which achieves this bound: the SCS auction that is 4-competitive.

6.2 Sampling Cost-Sharing Auction

Now we present another competitive truthful auction based on sampling that is simple, easy to analyze, and achieves the competitive ratio of four.

We first review a standard cost-sharing technique [10, 12]. The goal of this technique is, given bids $\mathbf{b}$ and cost C, to find a subset of the bidders to share the cost C. More precisely, the cost-sharing mechanism is defined as follows:

> CostShare_C: Given bids $\mathbf{b}$, find the largest k such that the highest k bidders' values are at lease C/k. Charge each C/k.

The two important properties of this mechanism are that

- CostShare_C is truthful.
- If $C \leq \mathcal{F}(\mathbf{b})$, CostShare_C has revenue C; otherwise it has no profit.

Because $\text{CostShare}_{\mathcal{F}'}$ and $\text{CostShare}_{\mathcal{F}''}$ are truthful on their respective partitions, we have,

Auction 4 Sampling Cost Sharing Auction (SCS):

1. Partition bids $\mathbf{b}$ uniformly at random into two sets, resulting in bid vectors $\mathbf{b}'$ and $\mathbf{b}''$.
2. Compute $\mathcal{F}' = \mathcal{F}(\mathbf{b}')$ and $\mathcal{F}'' = \mathcal{F}(\mathbf{b}'')$, the optimal fixed price profits for $\mathbf{b}'$ and $\mathbf{b}''$, respectively.
3. Compute the auction results by running CostShare$_{\mathcal{F}''}$ on $\mathbf{b}'$ and CostShare$_{\mathcal{F}'}$ on $\mathbf{b}''$.

Observation 6.6 *SCS is truthful.*

Next we show that SCS is competitive.

Theorem 6.7 *SCS is 4-competitive, and this bound is tight.*

Proof. In the special case that $\mathcal{F}' = \mathcal{F}''$ the auction profit is $\mathcal{F}' + \mathcal{F}'' \geq \mathcal{F}(\mathbf{b})$ and we are done. Otherwise, the auction profit is $R = \min(\mathcal{F}', \mathcal{F}'')$. Suppose, without loss of generality, that $\mathcal{F}' < \mathcal{F}''$. Then CostShare$_{\mathcal{F}''}$ on $\mathbf{b}'$ will reject all bids in $\mathbf{b}'$. However, CostShare$_{\mathcal{F}'}$ on $\mathbf{b}''$ will be able to achieve profit $\mathcal{F}'$.

By definition, $\mathcal{F}^{(2)}$ on $\mathbf{b}$ sells to $k \geq 2$ bidders at price p for a profit of $\mathcal{F}^{(2)} = kp$. These k bidders, all with bid value at least p, are divided uniformly at random between $\mathbf{b}'$ and $\mathbf{b}''$. Let k' be the number of them in $\mathbf{b}'$ and k'' the number in $\mathbf{b}''$. As such, $\mathcal{F}(\mathbf{b}') \geq pk'$ and $\mathcal{F}(\mathbf{b}'') \geq pk''$. Therefore,

$$\frac{\min(\mathcal{F}(\mathbf{b}'), \mathcal{F}(\mathbf{b}''))}{\mathcal{F}^{(2)}(\mathbf{b})} \leq \frac{\min(pk', pk'')}{pk} = \frac{\min(k', k'')}{k}.$$

Thus, the competitive ratio

$$\frac{\mathbf{E}[R]}{\mathcal{F}^{(2)}} = \tfrac{1}{k} \sum_{i=1}^{k-1} \min(i, k-i) \tbinom{k}{i} 2^{-k} = \frac{1}{2} - \tbinom{k-1}{\lfloor \frac{k}{2} \rfloor} 2^{-k}.$$

This ratio achieves its minimum of $1/4$ for $k = 2$ and $k = 3$. As k increases, the ratio approaches $1/2$.

To see that the bound presented on the competitive ratio is tight, consider the case where $\mathbf{b}$ consists of two very high bids h and $h+\epsilon$, and all other bids are negligibly small. In this case $\mathcal{F} = \mathcal{F}^{(2)} = 2h$, whereas the expected profit of the SCS auction is $h \cdot \mathbf{Pr}[\text{two high bids are split between } \mathbf{b}' \text{ and } \mathbf{b}''] = h/2 = \mathcal{F}/4$. ∎

The SCS auction, as described, gets no profit from one of the partitions, and therefore it is likely to lose at least half of the potential profit. An alternative is to pick a parameter $r < 1$ and run CostShare$_{(rF')}$ and CostShare$_{(rF'')}$. The competitive ratio of the resulting auction is $4/r > 4$. However, in a setting in which there is some prior knowledge of bidder distributions, a proper choice of r may lead to better performance.

7 Bounded Supply

Up to this point, we have studied the unlimited supply case, motivated by the digital goods market. In this section, we consider the case where the number of items available for sale is bounded. This case is typical for physical goods markets. We denote the number of items available by k. As before, the seller wishes to maximize profit and is not required to sell all the items. The definitions of truthful and competitive auctions, stated for the unlimited supply case, also apply to the bounded supply case. We denote by $\mathcal{F}^{(m,k)}$ the profit for the optimal single price auction that sells at least m and at most k items. It is this quantity that we wish to be competitive with.

To reduce the bounded supply case to the unlimited supply case, we can simply ignore (reject) all but the highest k bidders and run the unlimited supply auction on the remaining bids. We note that in order for this to be truthful we need to make sure that none of the bidders win at a price lower than the highest ignored bid. More formally the bounded supply auction $\mathcal{A}_k$ works as follows. Let $\mathbf{x}^{\mathcal{V}}$ and $\mathbf{p}^{\mathcal{V}}$ be the outcome of simulating k-Vickrey auction on $\mathbf{b}$. Let $\mathbf{x}^{\mathcal{A}}$ and $\mathbf{p}^{\mathcal{A}}$ be the outcome of simulating $\mathcal{A}$ on $\mathbf{b}^{\mathcal{A}}$ given by $b_i^{\mathcal{A}} = x_i^{\mathcal{V}} b_i$ (i.e., $\mathbf{b}$ with losers of k-Vickrey treated as zero). Compute the outcome of $\mathcal{A}_k$ as $\mathbf{p}$ with $p_i = \max(p_i^{\mathcal{V}}, p_i^{\mathcal{A}})$ and $\mathbf{x}$ with $x_i = x_i^{\mathcal{V}} x_i^{\mathcal{A}}$.

This technique allows us to trivially extend all results in this paper to the bounded supply case. Thus, for example, for DSOT we obtain:

Theorem 7.1 *The limited supply version of* DSOT *is constant competitive against* $\mathcal{F}^{(2,k)}$.

For SCS we obtain:

Theorem 7.2 *The limited supply version of SCS is 4-competitive against* $\mathcal{F}^{(2,k)}$.

It is interesting to point out that by using auctions like DSOT and SCS in a k-item auction it is possible on many bid vectors to obtain profit significantly higher than that of the k-item Vickrey auction.

8 Better than $\mathcal{F}$?

Thus far we have demonstrated several auction mechanisms that perform comparably to $\mathcal{F}^{(2)}$. In that $\mathcal{F}^{(2)}$ is comparable to $\mathcal{F}$, these auctions perform comparably to $\mathcal{F}$. As $\mathcal{F}$ is the optimal single-price auction it is clear that no auction that always uses a single price can achieve a higher revenue than $\mathcal{F}$. This gives good justification for comparing the revenue of single-price auctions to $\mathcal{F}$ (or $\mathcal{F}^{(2)}$).

The question remains of whether or not we have chosen the "best" metric possible for comparison. There are two primary goals in choosing the metric. First, we wish to achieve the best possible performance and thus we would like to compare ourselves against the strongest possible benchmark. On the other hand, we would like to find a natural metric that comes as close as possible to capturing the performance of the best truthful auction across a wide range of inputs.

Let us reconsider for a moment the strongest possible benchmark, the optimal omniscient multi-price auction, $\mathcal{T}$, that achieves as revenue the sum of the bids. Lemma 3.4 showed that the profit of $\mathcal{T}$ can be a factor of $\Theta(\log n)$ more than $\mathcal{F}$. There remains the question: Does there exist an alternative metric on bids $\mathbf{b}$ that captures the benefit of using multiple prices (and hence is sometimes larger than the $\mathcal{F}(\mathbf{b})$) that we could compare our truthful multi-priced auctions against?

As evidence to the contrary, we will show that no *monotone* auction can achieve an expected profit higher than $\mathcal{F}$ *on any input.* Monotone auctions are a large class of natural multi-priced truthful auctions; all the competitive auctions presented in this paper are monotone. This result supports the conjecture that there is no systematic way for an auction to achieve a higher profit than $\mathcal{F}$.

8.1 Hard-coded Auctions

We begin with some motivation for the notion of monotonicity by presenting some examples of non-monotone auctions. These auctions are not natural in the context of worst-case analysis in that while they can achieve significantly higher profit than that of $\mathcal{F}$ on certain bid sets (in fact they can achieve $\mathcal{T}(\mathbf{b})$), they do so at the cost of having horrendously low profit (and hence very bad competitive ratio) on other bid sets.

We first observe that for every $\mathbf{b}$ there is a symmetric truthful auction that achieves a profit of $\mathcal{T}(\mathbf{b})$. For example, consider the n-tuple $\mathbf{b}$ with half of the bids at value one and half at value $h > 2$. Thus, $\mathcal{T}(\mathbf{b}) = (h+1)n/2$ and $\mathcal{F}(\mathbf{b}) = hn/2$. Consider the symmetric auction given by bid-independent function f:

$$f(\mathbf{b}_{-i}) = \begin{cases} 1 & \text{if more } h\text{s than 1s in } \mathbf{b}_{-i}. \\ h & \text{otherwise.} \end{cases}$$

This auction achieves profit $\mathcal{T}(\mathbf{b})$ on our particular input: If $b_i = h$ then f outputs h; on $b_i = 1$, f outputs 1. Note that on most other inputs, $\mathbf{b}'$, this auction performs much worse than $\mathcal{F}$.

Now we generalize this result and show that for any set of bids $\mathbf{b}^*$ there exists a truthful (symmetric) auction that achieves a revenue of $\mathcal{T}(\mathbf{b}) = \sum_i b_i^*$. This is exemplified in the bid-independent auction $\mathcal{A}_{f_{\mathbf{b}^*}}$ parameterized by $\mathbf{b}^*$ and defined as follows:

$$f_{\mathbf{b}^*}(\mathbf{b}_{-i}) = \begin{cases} b_j & \text{if } \pi(\mathbf{b}_{-i}) = \mathbf{b}^*_{-j} \text{ for some permutation } \pi \\ \infty & \text{otherwise.} \end{cases}$$

The "otherwise" case is arbitrarily chosen. In fact any number of bid vectors that have a pairwise difference of at least two bid values can be hard-coded into an auction in this manner. The auction will perform very poorly on any input that differs on only one bid value from one of the hard-coded bid vectors.

For worst case profit maximization, the mechanism of both of these auctions is counter intuitive. For the case that half the bids are at one and half at h, the bid-independent function sees more h values and outputs one. When it sees less h values, it outputs h. A more intuitive output for a profit maximizing auction would be to output h when there are more bids at value h.

Note that we can combine a hard-coded auction and a competitive auction by flipping a fair coin and running the former or the latter depending on the outcome of the toss. The resulting auction is competitive, with the competitive ratio twice that of the underlying competitive auction. Furthermore, on the hard-coded input $\mathbf{b}^*$, the expected revenue of the auction is at least $\mathcal{T}(\mathbf{b}^*)/2$, which can be significantly bigger than $\mathcal{F}(\mathbf{b}^*)$.

Although a competitive auction can outperform $\mathcal{F}$ on some inputs, we conjecture that this happens at the expense of the competitive ratio and auctions designed to achieve high competitive ratios do not outperform $\mathcal{F}$ in this sense. In the next section we introduce the class of monotone auctions that includes DSOT, SCS, DOT, and Vickrey with reservation price. We show that no monotone auction can outperform $\mathcal{F}$.

8.2 Monotonicity

The intuition underlying our notion of monotonicity is that the bid-independent function defining the auction should output higher prices when it sees higher bid values.

Definition 8.1 *An auction is* monotone *if for any pair of bidders i and j with $b_i \leq b_j$, we have:*

$$\forall x \leq b_i,\ \mathbf{Pr}[\textit{bidder } i \textit{ wins at price} \leq x] \leq \mathbf{Pr}[\textit{bidder } j \textit{ wins at price} \leq x]\,.$$

The intuition is that if $b_i < b_j$, then $\mathbf{b}_{-i}$ looks like a higher set of bids than $\mathbf{b}_{-j}$. Therefore, the price bidder i pays if he wins will tend to be higher than the price bidder j pays if he wins.

The class of monotone auctions is very general. It is not difficult to verify that the Vickrey auction with a reservation price is monotone. Thus, the optimal Bayesian auctions for i.i.d. prior distributions are monotone. Analysis of DSOT, SCS, and DOT, the auctions introduced in this paper, shows that they are also monotone. The proofs of these facts are given in Appendix B. Our optimal single-price omniscient auction, $\mathcal{F}$, is also monotone.

We now show that $\mathcal{F}$ is the optimal monotone auction. This further justifies our comparison to $\mathcal{F}$ and the related metric $\mathcal{F}^{(2)}$.

Theorem 8.2 *Let $\mathcal{A}$ be any monotone (truthful) randomized auction. For all bid vectors $\mathbf{b}$, the revenue $R = \sum_i p_i$ of $\mathcal{A}$ on input $\mathbf{b}$ satisfies*

$$\mathbf{E}[R] \leq \mathcal{F}(\mathbf{b}).$$

Proof. Let f be the bid-independent function defining $\mathcal{A}$. For each i, define $g_i(x) = \mathbf{Pr}[f(b_{-i}) \leq x]$.

Now consider the following thought experiment. Let U be a random variable that is uniform on $[0, 1]$. Imagine running the bid-independent auction that for each i uses $g_i^{-1}(U)$ to set the threshold for bidder i, with g_i^{-1} defined as $g_i^{-1}(y) = \inf\{x\ :\ g_i(x) = y\}$. We denote by R_U the resulting auction revenue. We observe that the threshold distribution for bidder i in this experiment is precisely the same as the original threshold distribution for bidder i:

$$\mathbf{Pr}\left[g_i^{-1}(U) \leq x\right] = \mathbf{Pr}[U \leq g_i(x)] = g_i(x).$$

Therefore, by summing the expectations for the bidders, we obtain

$$\mathbf{E}[R_U] = \mathbf{E}[R] \,.$$

We complete the proof by showing that the expected revenue from our thought experiment $\mathbf{E}[R_U]$ is at most $\mathcal{F}(\mathbf{b})$. Conditioned on $U = u$, let k be the index of the smallest winning bid. Thus, $g_k^{-1}(u) \leq b_k$. Since $\mathcal{A}$ is monotone, for $x \leq b_k$ and all j with $b_j \geq b_k$, we have $g_k(x) \leq g_j(x)$. Furthermore, $g_k(x)$ and $g_j(x)$ are monotone non-decreasing functions. Therefore, it must be that $g_j^{-1}(u) \leq g_k^{-1}(u) \leq b_k \leq b_j$ and therefore all bidders with bid values at least b_k win at a price at most b_k. Thus, the revenue, R_u, is at most b_k times the number of bids with bid value least b_k which totals to at most $\mathcal{F}(\mathbf{b})$. This holds for all $u \in [0, 1]$, and thus $\mathbf{E}[R_U] \leq \mathcal{F}(\mathbf{b})$. ∎

9 Concluding Remarks

In this paper, we introduced a framework for designing and analyzing profit-maximizing truthful auctions that are competitive against optimal auctions on all inputs. We provided motivation for the framework and presented two competitive auctions.

A number of interesting open problems remain. In this paper we have shown an upper bound of 4 and a lower bound of 2.42 on the competitive ratio against $\mathcal{F}^{(2)}$. It would be interesting to bridge the gap. More generally, we would like to understand the precise tradeoff between m and β for auctions that are β-competitive against $\mathcal{F}^{(m)}$.

The parameterized DSOT$_r$ family of auctions can achieve competitive ratios which approach 1 as m increases. A more elegant result would be a single, non-parameterized auction that has the same property.

An interesting question is how to extend the competitive framework introduced in this paper to other mechanism design problems. A followup paper [5] makes progress in this direction by introducing the concept of a cancellable auction, a competitive auction that can be cancelled if its revenue fails to meet a target. The paper shows that the DSOT auction becomes untruthful if cancelling is allowed while the SCS auction remains truthful, i.e., it is *cancellable*. Cancellable auctions apply to a wider range of problems and can be composed to build more general mechanisms.

Finally, as we have observed several times in this paper, competitive (worst-case) analysis and the traditional (Bayesian) analysis of auctions are not mutually exclusive. An interesting direction for research is Bayesian analysis of our auctions and their variants. We conjecture that such analysis can lead to much tighter bounds than the worst-case bounds we present.

Acknowledgments

We would like to thank the following people for their contributions and useful discussions relating to this work: Amos Fiat, Stuart Haber, Jim Horning, Umesh Maheshwari, Satish Rao, Amir Ronen, Ilya Segal, Eva Tardos, Bob Tarjan, Radek Vingralek, and Stephen Weeks. We would also like to thank the anonymous referees for their help in improving the presentation.

References

[1] N. Abel. Beweis eines Ausdrucks von welchem die Binomial-Formel ein einzelner Fall ist. *Crelles Journal für die Reine und Angewandte Mathematik*, 1:159–160, 1826.

[2] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[3] J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.

[4] E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.

[5] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. 34th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2002.

[6] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. Technical Report STAR-TR-99.09.01, STAR Laboratory, InterTrust Tech. Corp., Santa Clara, CA, 1999. Available at URL http://www.star-lab.com/tr/tr-99-01.html.

[7] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Alg.*, pages 735–744. ACM/SIAM, 2001.

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

[9] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[10] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs. *Economic Theory*, to appear.

[11] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.

[12] L. S. Shapley. Core of Convex Games. *Int. J. of Game Theory*, 1:11–26, 1971.

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

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

A Analysis of DSOT

Below we prove that DSOT is competitive. Moreover, we show that, as m gets large, the expected profit of DSOT_r tends to $\mathcal{F}^{(m)}$ for $r = \lceil \frac{m}{2}(1-\delta) \rceil$. This implies that for any $\epsilon > 0$, we there is m' such that for any $m \geq m'$, DSOT_r is $(1+\epsilon)$-competitive against $\mathcal{F}^{(m)}$.

A.1 Preliminaries

We begin our analysis of DSOT_r with some preliminary lemmas. We will use the following definitions:

- For any value v, let $F_v = vn_v$ (resp. $F'_v = vn'_v$ and $F''_v = vn''_v$) where n_v is the number of bids in $\mathbf{b}$ (resp. $\mathbf{b}'$ and $\mathbf{b}''$) greater than or equal to v. Thus, F_v is the profit from using v as a threshold for $\mathbf{b}$.

- Let R be the random variable representing the profit of the auction DSOT_r (where the input $\mathbf{b}$ is implied by the context). Recall from the definition of the DSOT_r auction that $p' = \mathrm{opt}_r(\mathbf{b}')$ and $p'' = \mathrm{opt}_r(\mathbf{b}'')$ are the thresholds used for $\mathbf{b}''$ and $\mathbf{b}'$ respectively. Thus, the profit of the auction is $R = F'_{p''} + F''_{p'}$.

- Let E_α be the event

$$E_\alpha: \quad R \geq (1-\alpha)\left(F_{p'} + F_{p''}\right). \tag{1}$$

  for $0 \leq \alpha \leq 1$. This event holding for small α indicates a high profit, a significant fraction of the total profit achievable using p' and p'' as thresholds.

- Let $r = \lceil \frac{m}{2}(1-\delta) \rceil$. Suppose that $\mathrm{opt}_m(\mathbf{b}) = v_k$, the k-th largest bid in $\mathbf{b}$. Define H to be the event that (a) there are at least r bids in $\mathbf{b}'$ that are at least v_k and (b) there are at least r bids in $\mathbf{b}''$ that are at least v_k. Note that if H does not occur then DSOT_r may be unable to pick a threshold of magnitude similar to v_k from at least one of the partitions.

The key lemma we use is the following:

Lemma A.1 *Let $0 < \delta < 1$ and $r = \lceil \frac{m}{2}(1-\delta) \rceil$. Then, for any $\mathbf{b}$, DSOT_r satisfies*

$$\mathbf{Pr}\left[R \geq \frac{1-\alpha}{\alpha}\mathcal{F}^{(m)}(\mathbf{b})\right] \geq \mathbf{Pr}[E_\alpha \cap H]$$

and thus

$$\mathbf{E}[R] \geq \frac{1-\alpha}{\alpha}\mathcal{F}^{(m)}(\mathbf{b})\mathbf{Pr}[E_\alpha \cap H]\,.$$

Proof. Suppose that $\mathcal{F}^{(m)}(\mathbf{b}) = kv_k$. By definition, $\mathrm{opt}_r(\mathbf{b}')$ returns the p' that maximizes $F'_{p'}$ conditioned on there being r bids that are at least p' (likewise for $\mathbf{b}''$ and p''). Notice that, conditioned on H, there are at least r elements in each of $\mathbf{b}'$ and $\mathbf{b}''$ that are at least v_k. Thus, we can conclude that

$$F'_{v_k} \leq F'_{p'} \quad \text{and} \quad F''_{v_k} \leq F''_{p''}.$$

We also have

$$\mathcal{F}^{(m)}(\mathbf{b}) = F_{v_k} = F'_{v_k} + F''_{v_k},$$

and thus

$$\mathcal{F}^{(m)}(\mathbf{b}) \leq F'_{p'} + F''_{p''}. \tag{2}$$

Event E_α holding, equation (1) allows us to lower bound the profit R as

$$\begin{aligned} R &\geq (1-\alpha)[F_{p'} + F_{p''}] \\ &= (1-\alpha)[F'_{p'} + F''_{p'} + F'_{p''} + F''_{p''}] \\ &= (1-\alpha)[R + F'_{p'} + F''_{p''}] \end{aligned}$$

which by equation (2) gives

$$R \geq (1-\alpha)[R + \mathcal{F}^{(m)}].$$

Event H holding, we rearrange terms to obtain

$$R \geq \tfrac{1-\alpha}{\alpha}\mathcal{F}^{(m)}.$$

■

This lemma reduces our problem to studying the events E_α and H.

To do so, for any j, we define B_j to be the j highest bids in $\mathbf{b}$ (i.e., $\{v_1, \ldots v_j\}$), and let $n'(B_j)$ be the number of these bids that are in $\mathbf{b}'$.

Definition A.2 *Given c: $0 < c < 1$, we say that B_j is c-*good *if*

$$\lceil cj \rceil \leq n'(B_j) \leq j - \lceil cj \rceil.$$

*Otherwise, B_j is c-*bad.

We prove that B_j is likely to be c-good using the following version of the Chernoff bound:

Theorem A.3 (see e.g. [9], page 70) *Let X_i, $1 \leq i \leq n$ be independent Bernoulli trials such that for all i, $\mathbf{Pr}[X_i = 1] = 1/2$. Then for $X = \sum_{1\leq i \leq n} X_i$, and $0 < \delta \leq 1$,*

$$\mathbf{Pr}\Big[X < (1-\delta)\frac{n}{2}\Big] < e^{-\frac{\delta^2 n}{4}}.$$

Since the partition of $\mathbf{b}$ into the two subvectors $\mathbf{b}'$ and $\mathbf{b}''$ is done by flipping a fair coin for each bid, we can conclude from the Chernoff bound that

$$\mathbf{Pr}[B_j \text{ is } c\text{-bad}] \leq 2e^{-\frac{(1-2c)^2 j}{4}}.$$

Thus, we can conclude by a simple union bound that

$$\mathbf{Pr}[B_j \text{ is not } c\text{-good for some } j > t\ ] \leq \sum_{j \geq t} 2e^{-\frac{(1-2c)^2 j}{4}}$$

and therefore we have

Lemma A.4

$$\mathbf{Pr}[B_j \textit{ is not c-good for some } j > t] \leq \frac{2e^{-\frac{(1-2c)^2 t}{4}}}{1 - e^{-(1-2c)^2/4}}.$$

A.2 Analysis of DSOT

We are now ready to proceed with the analysis of DSOT_r. First we show that as m increases, the competitive ratio of the auction approaches one.

Theorem A.5 *Let $\delta: \; 0 < \delta < 1$ be a constant and let $r = \lceil \frac{m}{2}(1-\delta) \rceil$. Then in the limit as $m \to \infty$, DSOT_r is $(1+\frac{\delta}{2})/(1-\frac{\delta}{2})$ competitive against $\mathcal{F}^{(m)}$.*

There is an absolute constant $C > 0$ such that

$$\mathbf{Pr}\left[\mathrm{DSOT}_r \geq \frac{(1-\frac{\delta}{2})}{(1+\frac{\delta}{2})}\mathcal{F}^{(m)}\right] \geq 1 - e^{-C\delta^2 m}.$$

Proof. Fix $\epsilon = \delta/2$. An immediate corollary of Lemma A.4 is that

$$\lim_{m\to\infty} \mathbf{Pr}\big[B_j \text{ is } \tfrac{1}{2}(1-\epsilon)\text{-good for all } j > \tfrac{m}{2}(1-2\epsilon)\big] = 1 - o(1). \tag{3}$$

Thus, with probability $1-o(1)$, $F'_{p'} \leq \frac{1}{2}(1+\epsilon)F_{p'}$ and $F''_{p''} \leq \frac{1}{2}(1+\epsilon)F_{p''}$, and thus

$$\mathbf{Pr}\Big[E_{\frac{1}{2}(1+\epsilon)}\Big] = 1 - o(1). \tag{4}$$

Finally, we show that $\mathbf{Pr}[H] = 1-o(1)$. As before, we assume that $\mathrm{opt}_m(\mathbf{b})$ is the k-th largest bid v_k. We also assume, without loss of generality, that v_k is in $\mathbf{b}'$, and that v_ℓ is the smallest bid larger than v_k that is in $\mathbf{b}''$. Let G be the event that there is at least one element v_i with $m(1-\epsilon) \leq i \leq m$ in each of $\mathbf{b}'$ and $\mathbf{b}''$. Then

$$\lim_{m\to\infty} \mathbf{Pr}[G] = \lim_{m\to\infty}\left(1 - 2^{1-\epsilon m}\right) = 1 - o(1).$$

Thus, with probability $1-o(1)$,

$$m(1-\epsilon) \leq \ell < k. \tag{5}$$

Moreover, from (3), we can conclude that with probability at least $1-o(1)$, there are at least $\frac{m}{2}(1-\epsilon) > r$ bids above v_k in $\mathbf{b}'$. Also, from (3) and (5), we can conclude that there are at least $\frac{m}{2}(1-\epsilon)(1-\epsilon) > r$ bids above v_ℓ in $\mathbf{b}''$. Thus,

$$\mathbf{Pr}[H] = 1 - o(1). \tag{6}$$

Finally, from Equations (4) and (6) and Lemma A.1, we have

$$\mathbf{Pr}\left[R \geq \frac{(1-\epsilon)}{(1+\epsilon)}\mathcal{F}^{(m)}\right] = 1 - o(1),$$

where R is the profit of DSOT_r. ■

Next we show that DSOT is β-competitive for a (relatively large) constant β.

Theorem A.6 *There is a constant β such that* DSOT *is β-competitive.*

Proof. As before, we denote by v_i the value of the i-th largest bid. Suppose that $opt_2(\mathbf{b})$ is the k-th largest bid. We restrict our attention only to partitions of the bids in which the highest bid v_1 is in one subvector, without loss of generality, in $\mathbf{b}'$, and both v_2 and v_k are in the other subvector. This event, which we shall denote G, occurs with probability $1/4$.

We claim that

$$\mathbf{Pr}\left[E_{\frac{19}{20}}\right] \geq 0.05. \tag{7}$$

Indeed, an immediate corollary of Lemma A.4 is that

$$\mathbf{Pr}\left[B_j \text{ is } \tfrac{1}{20}\text{-good for all } j > 20\right] \geq 0.8.$$

Moreover, if event G holds, all $j \leq 20$ are $\frac{1}{20}$-good, since v_1 is in $\mathbf{b}'$ and v_2 is in $\mathbf{b}''$. Thus, with probability at least $0.8 - \mathbf{Pr}[\neg G] = 0.8 - 0.75 = 0.05$, all B_k are $\frac{1}{20}$-good and event G holds, and thus, in particular, event $E_{\frac{19}{20}}$ holds.

For the case $m = 2$, event G implies event H. Thus, from (7) and Lemma A.1,

$$\mathbf{Pr}\left[R \geq \frac{1}{19}\mathcal{F}^{(2)}\right] \geq 0.05,$$

and thus DSOT is 380-competitive.

■

B Monotonicity of DOT, DSOT, and SCS

B.1 DOT is monotone

As DOT is a symmetric auction, the order of the bids in the input does not affect the outcome. For convenience of notation we will assume that they are indexed from highest, b_1, to lowest, b_n. We denote by $t = \text{opt}(\mathbf{b})$ and $t_1 = \text{opt}(\mathbf{b}_{-1})$ the optimal sale price for $\mathbf{b}$ and $\mathbf{b}_{-1}$, respectively.

Theorem B.1 *DOT is monotone as*

- *It uses the single sale price t_1 for all winners.*
- *All bidders that bid above the lowest winning bid also win.*

Proof. Recall that the bid-independent function that implements DOT is $f(\mathbf{b}) = \text{opt}(\mathbf{b}) = \text{argmax}_{b_k} kb_k$. We will show that for any bidder i that wins the auction, bidder $i-1$ also wins the auction and at the same price. From this, a simple induction gives the theorem.

In the case that $b_i = b_{i-1}$ the fact that DOT is symmetric implies that they must both win at the same price. Now consider the case that $b_{i-1} > b_i$. Since bidder i wins, the computation of $\text{opt}(\mathbf{b}_{-i})$ must find the maximum of kb_k for $k < i$ and $(k-1)b_k$ for $k > i$ to be $(k^*-1)b_{k^*}$ for $k^* > i$. The only difference between the computation of $\text{opt}(\mathbf{b}_{-i})$ and of $\text{opt}(\mathbf{b}_{-i+1})$ is that for $\text{opt}(\mathbf{b}_{-i})$ we consider $(i-1)b_{i-1}$ and for $\text{opt}(\mathbf{b}_{-i+1})$ we consider $(i-1)b_i$. Since $\text{opt}(\mathbf{b}_{-i})$ finds $(k^*-1)b_{k^*}$ bigger than the other values it considers and since $(i-1)b_{i-1} > (i-1)b_i$, it must be that $(k^*-1)b_{k^*}$ is the biggest value that $\text{opt}(\mathbf{b}_{-i+1})$ considers as well. ■

B.2 DSOT Auction

We will show that DSOT is monotone with respect to any two bidders i and j. First fix i and j such that $b_i \leq b_j$. The randomness of DSOT is in the partitioning. We verify the monotonicity of DSOT by presenting a disjoint grouping of the partitionings such that each group is monotone with respect to i and j and the DSOT auction is just a randomization over monotone groups.

We put partitionings with b_i and b_j in the same partition in their own group. Note that these groups are monotone with respect to b_i and b_j because both b_i and b_j are offered the same price.

Any partitioning with b_i and b_j in opposite partitions we will pair in a group with the partitioning we get when we swap b_i with b_j. This gives us a group with two partitionings:

Partitioning P: $B_1 \cup \{b_i\}$ and $B_2 \cup \{b_j\}$.

Partitioning P': $B_1 \cup \{b_j\}$ and $B_2 \cup \{b_i\}$.

We show that this group $\{P, P'\}$ is monotone with respect to b_i and b_j if one of the two partitionings is chosen by flipping a fair coin. Consider the thought experiment where for b_i a "heads" coin flip means use P and "tails" means use P', but for b_j a "heads" coin flip means use P' and "tails" for P. In this experiment, "heads" results in b_i getting the optimal price for $B_2 \cup \{b_j\}$ and b_j getting the optimal price of $B_2 \cup \{b_i\}$. This is just the outcome of DOT on the bid set $B_2 \cup \{b_i, b_j\}$. Likewise on "tails", the outcome is that of DOT on the bid set $B_1 \cup \{b_i, b_j\}$. By Theorem B.1, both of these outcomes are monotone. Note that the outcome for b_i and b_j as random variables in our thought experiment are the same as the random variables for their actual outcomes for this grouping. Thus, this grouping is monotone with respect to bidders i and j.

B.3 SCS Auction

To show that SCS is monotone, we will use the same general approach as for DSOT of finding disjoint groupings in the partitionings and showing that each is itself monotone with respect to two bids b_i and b_j with $b_i \leq b_j$.

As with DSOT, partitionings with both b_i and b_j on the same partition are themselves monotone as the cost sharing mechanism is monotone: it gives an outcome such that all winning bidders pay the same price and all bidders whose bid value is above this price win.

Likewise, we pair a partitioning with b_i and b_j in different partitions, with the partitioning with b_i and b_j swapped. We get partitionings P and P' as defined above.

$$F_j = \mathcal{F}(B_2 \cup \{b_j\}) \qquad F_i = \mathcal{F}(B_1 \cup \{b_i\})$$
$$F'_i = \mathcal{F}(B_2 \cup \{b_i\}) \qquad F'_j = \mathcal{F}(B_1 \cup \{b_j\})$$

Since $b_i \leq b_j$ we have:

$$F_j \geq F'_i \qquad F'_j \geq F_i$$

Note that if b_i does not win for either P or P' then this grouping is trivially monotone with respect to bidders i and j. Otherwise, suppose b_i wins in partitioning P and pays price p_i. We will show that b_j wins in P' and pays $p'_j \leq p_i$.

Let cs_C be the bid-independent function for $\mathrm{CostShare}_C$. Recall that for partitioning P, the price for b_i is computed by running $\mathrm{CostShare}_{F_j}$ on $B_1 \cup \{b_i\}$. For partitioning P' the price for b_j is computed by running $\mathrm{CostShare}_{F'_i}$ on $B_1 \cup \{b_j\}$. Thus,

$$p_i = \mathrm{cs}_{F_j}(B_1) \geq \mathrm{cs}_{F'_i}(B_1) = p_j.$$

The intermediate step here follows because $F_j \geq F'_i$ and because cs_C is monotone in C.

C Technical Proof for Section 5

First we prove Lemma 5.2 that says that for the prior distribution defined in Section 5, we have

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

Proof. To get a bound on the expected value of $\mathcal{F}^{(2)}$ on n bids from our distribution we define a recurrence based on $F_{n,k}$ defined as

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

for bids $\mathbf{b}^{(n)}$ sorted from highest to lowest (i.e., $b_i \geq b_{i+1}$). To define the recurrence, fix n, k, and z and consider the events $\mathcal{G}_i$ and $\mathcal{H}_i$ for $1 \leq i \leq n$ defined as

$$\begin{aligned}\mathcal{G}_i &= i \text{ of the } n \text{ bids are at least } (k+i)/z\\ \mathcal{H}_i &= \mathcal{G}_i \wedge \neg \bigvee_{j=i+1}^{n} \mathcal{G}_j.\end{aligned}$$

We have

$$\begin{aligned}\mathbf{Pr}[\mathcal{G}_i] &= \binom{n}{i}\left(\frac{k+i}{z}\right)^i\\ \mathbf{Pr}[\mathcal{H}_i] &= \mathbf{Pr}[\mathcal{G}_i]\,\mathbf{Pr}\left[\neg\bigvee_{j=i+1}^{n}\mathcal{G}_j \;\middle|\; \mathcal{G}_i\right].\\ &= \mathbf{Pr}[\mathcal{G}_i]\,\mathbf{Pr}[F_{n-i,k+i} < z]\\ &= \binom{n}{i}\left(\frac{k+i}{z}\right)^i\mathbf{Pr}[F_{n-i,k+i} < z]\end{aligned}$$

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,

$$
\begin{aligned}
\mathbf{Pr}[F_{n,k} > z] &= \mathbf{Pr}\left[\bigwedge_{i=1}^{n} \mathcal{G}_i\right] = \sum_{i=1}^{n} \mathbf{Pr}[\mathcal{H}_i] \\
&= \sum_{i=1}^{n} \binom{n}{i} \left(\frac{k+i}{z}\right)^i \mathbf{Pr}[F_{n-i,k+i} < z] \qquad (8)
\end{aligned}
$$

Also, note that $F_{0,k} = 0$. For n bids $\mathbf{b}^{(n)}$, $\mathcal{F}(\mathbf{b}^{(n)}) = F_{n,0}$. We are interested in $\mathcal{F}^{(2)}(\mathbf{b}^{(n)})$ which is the same as $\mathcal{F}(\mathbf{b}^{(n)}) = F_{n,0}$ except that we ignore the $\mathcal{H}_1$ case. This gives

$$
\begin{aligned}
\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] &= \mathbf{Pr}[F_{n,0} > z] - \mathbf{Pr}[\mathcal{H}_1] \\
&= \mathbf{Pr}[F_{n,0} > z] - \tfrac{n}{z}\mathbf{Pr}[F_{n-1,1} < z]\,. \qquad (9)
\end{aligned}
$$

So in order to obtain $\mathbf{Pr}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})\big]$ we need to solve the recurrence $F_{n,k}$, i.e., Equation (8). We will show that the solution is:

$$
\mathbf{Pr}[F_{n,k} > z] = 1 - \left(\frac{z-k}{z}\right)^n \left(\frac{z-k-n}{z-k}\right). \qquad (10)
$$

Note that our solution for the recurrence is correct for $n = 0$. We show that it is true in general inductively.

$$
\mathbf{Pr}[F_{n,k} > z] = \sum_{i=1}^{n} \binom{n}{i} \left(\frac{k+i}{z}\right)^i \mathbf{Pr}[F_{n-i,k+i} < z]\,.
$$

Substituting in our solution, we get

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

We now apply the following version of Abel's Identity [1]

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

Making the change of variables, $j = n - i$, $z = -1$, $x = z - k - n$, and $y = k + n$ we get:

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

We plug this in above and subtract out the $i = 0$ term to get

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

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

Recall our goal is to compute $\mathbf{Pr}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\big]$. Equation (10) shows that $\mathbf{Pr}[F_{n,0} > z] = n/z$. This combined with Equation (9) and Equation (10) gives

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

To complete this proof, we use the formula $\mathbf{E}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})\big] = \int_0^\infty \mathbf{Pr}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\big]\, dz = n + \int_n^\infty \mathbf{Pr}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\big]\, dz$. In the form above, this is not easily integrable, However, we can transform it back into a binomial sum which we can integrate:

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

■

Next we prove Lemma 5.4 that claims

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

Proof. 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 use the following fact below: If, for $1 \le k \le K$, $0 < a_k < 1$, then

$$\prod_{k=1}^{K}(1 - a_k) \ge 1 - \sum_{k=1}^{K} a_k.$$

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

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