Competitive Generalized Auctions

Amos Fiat
School of Comp. Sci.
Tel Aviv University
Tel Aviv, Israel
fiat@tau.ac.il

Andrew V. Goldberg
STAR Lab
InterTrust Tech. Corp.
Santa Clara, CA 95054-1851
avg@acm.org

Jason D. Hartline*
Dept. Comp. Sci.
University of Washington
Seattle, WA 98195-2350
hartline@cs.washington.edu

Anna R. Karlin†
Dept. Comp. Sci.
University of Washington
Seattle, WA 98195-2350
karlin@cs.washington.edu

God help us if we ever take the theater out of the auction business or anything else. It would be an awfully boring world.

— *A. Alfred Taubman,*
Chairman, Sotheby Galleries.
Wall Street Journal, 18 Sep 1985.

Perhaps no one as yet has been truthful enough about what "truthfulness" is.

— *Friedrich Nietzsche.*
Beyond Good and Evil, (1886).

ABSTRACT

We describe mechanisms for auctions that are simultaneously truthful (alternately known as strategy-proof or incentive-compatible) and guarantee high "net" profit. We make use of appropriate variants of competitive analysis of algorithms in designing and analyzing our mechanisms. Thus, we do not require any probabilistic assumptions on bids.

We present two new concepts regarding auctions, that of a cancellable auction and that of a generalized auction. We use cancellable auctions in the design of generalized auctions, but they are of independent interest as well. Cancellable auctions have the property that if the revenue collected does not meet certain predetermined criteria, then the auction can be cancelled and the resulting auction is still truthful. The trivial approach (run a truthful auction and cancel if needed) yields an auction that is not necessarily truthful.

Generalized auctions can be used to model many problems previously considered in the literature, as well as numerous new problems. In particular, we give the first truthful profit-maximizing auctions for problems such as conditional financing and multicast.

*Research supported in part by NSF grant CCR-0105406.
†Research supported in part by NSF grant CCR-0105406.

STOC'02, May 19-21, 2002, Montreal, Quebec, Canada.

1. INTRODUCTION

Recent economic and computational trends, such as the negligible cost of duplicating digital goods and the emergence of the Internet as one of the most important arenas for resource sharing between parties with diverse and selfish interests, have led to a number of new and interesting dynamic pricing problems. In this paper, we describe a generalized auction framework within which many of these problems can be formulated and solved. Our focus is on the design of auction mechanisms that maximize the profit of the resource provider (a.k.a., the auctioneer), assuming selfish behavior on the part of the resource consumers.[1]

We follow a common approach originating in the field of *mechanism design*, which is to design protocols in which all rational participants are motivated to behave according to the protocol. We assume that each consumer has a private *utility value*, the maximum value that the consumer is willing to pay for the resource. For dynamic pricing problems, the *mechanism* takes as input *bids* from each of the consumers and determines which bidders receive the resource and at what price. We say the mechanism is *truthful* (or equivalently, strategy-proof or incentive-compatible) if it is in each consumer's best interest to bid their true utility value.

1.1 Motivating Problems

Basic Auction: An auctioneer has k identical items to be sold in a sealed-bid, single-round auction. This problem is, of course, well-studied by economists and game theorists (see, e.g., the survey on auction theory by Klemperer [11]). The traditional model of bidder behavior is that bidders try to maximize their profit, defined as the difference between their utility value and the price at which they get the item. Given this definition, the only truthful mechanism which sells the k items to the top k bidders is the classical k-item Vickrey auction [21], in which each bidder pays the $k + 1$st

[1] In our discussion below, we shall interchangeably refer to the notion of providing a good or a service or a resource.

bid value.

A natural problem is whether or not there are truthful mechanisms whereby the auctioneer can increase their revenue by selling fewer items.[2] For example, it could be that selling only 1 item, say using a 1-item Vickrey auction, the auctioneer could obtain a much greater revenue than by selling all k items. This would happen, for example, if the second highest bid is at least k times the $k+1$st highest bid value. The question then becomes: how do we design truthful auctions that obtain near-optimal revenue to the auctioneer, for any values of the bidders' bids?

Goldberg, Hartline, and Wright [7] looked at this problem for the case where k is unbounded and defined a *competitive framework* for its analysis. This competitive framework allows auctions to be analyzed without Bayesian assumptions on consumer utility values. They also presented an auction that is guaranteed to get near-optimal revenue for the case where there is sufficient competition between the bidders (meaning that an optimal mechanism would sell many items). The problem of designing a revenue maximizing truthful auction when the competition for the item is minimal was left open in their work.

Conditional Financing: A company is considering an initial public offering of its shares, or a venture capitalist is trying to sell fixed return junk bonds for some venture. The company would like to sell the shares only if the revenue raised is sufficiently high. How can this be done so that the company raises as much money as is possible, given the potential shareholders' honest valuations of the company, but so that, if the company is not valued sufficiently highly by the public, the IPO can be cancelled?

Offering a pay-per-view broadcast in a segmented market: A company offering a pay-per-view broadcast needs to formulate a pricing scheme for its potential viewers. Suppose that the potential viewers are partitioned into markets (e.g., by location or by some measure of the quality of the good they are receiving). Further, suppose that the cost of providing the broadcast to viewers in the ith market is C_i, a fixed cost which is paid once if and only if there are any viewers in the ith market. The goal of the company is to maximize its profit, the sum of the prices paid by each of the viewers minus the costs of providing the broadcast to those markets in which there are viewers.

Multicast pricing: Feigenbaum, Papadimitriou and Shenker [5] initiated the study of pricing algorithms for multicast transmission. The model is a network with users residing at nodes in the network. There are costs associated with transmitting data across each of the links in the network. Each user has a utility for receiving the broadcast. The problem is to choose the multicast tree and the prices to charge each of the recipients of the broadcast. Feigenbaum et al. focus on the network complexity of implementing solutions that are either *budget-balanced*, in which the broadcaster precisely recovers the cost of the transmission, or *efficient*, in which the receivers chosen are the set which maximizes the difference between the sum of the utilities of the receivers and the cost of multicasting to that set of receivers. For budget-balanced solutions there is no profit to the broadcaster, and for efficient solutions the broadcaster may run a deficit. Our interest is in designing truthful pricing mechanisms which maximize the broadcaster's profit. This natural problem has received no attention to date.

[2]Surprisingly, to our knowledge, this problem has not been considered by economists. Most likely, this is because their study of auctions typically uses a Bayesian approach that assumes the knowledge of probability distributions on bidder utility values.

1.2 Generalized Auction Problems

We define a framework for modeling dynamic pricing problems for revenue maximization as *generalized auctions*. This framework captures all of the problems just described, and many others.

A *generalized auction problem* $\mathcal{A} = (\mathcal{S}, c(\cdot))$ is described by the following parameters:

- $\mathcal{S} = \{S_i : 1 \le i \le m\}$. The sets S_i partition the n bidders into m *market segments*. Importantly, within each market, bidders are indistinguishable from one another.[3]

- $c(\cdot)$, a cost function mapping vectors in $\{0,1\}^m$ to non-negative real numbers. The domain of $c(\cdot)$ describes each possible market allocation, so, for example, a vector $\mathbf{r} = (r_1, \ldots, r_m)$ in the domain indicates for each market i whether or not goods are allocated to that market ($r_i = 1$ or $r_i = 0$, respectively). The cost, $c(\mathbf{r})$, is the cost to the auctioneer (the service provider) of providing the goods assuming the market allocation, $\mathbf{r}$.

So, for example, for the basic unlimited supply auction problem, $m = 1$, and $c(\mathbf{r}) = 0$ for all $\mathbf{r}$ as duplicating and distributing the goods is assumed to be free for the auctioneer. For the multicast pricing problem, the number of markets is the number of nodes in the network, and the viewers at the ith node form S_i, the ith market. The cost of a market allocation $c(\mathbf{r})$ is exactly the cost of transmitting along the multicast tree defined by the nodes (markets) receiving the broadcast.

A *mechanism* for a *generalized auction problem* takes as input a bid value b_i for each of the n bidders (which we combine to get the bid vector $\mathbf{b}$), and decides for each bidder, whether or not that bidder receives the good and, if so, at what price. The *profit* of the auctioneer is the difference between the prices paid by the receiving bidders and the cost of providing the goods to those bidders. Our goal is to design mechanisms that are truthful and which maximize the profit of the auctioneer.

Following [7], we use a form of competitive analysis to evaluate the performance of truthful mechanisms. Consider a particular generalized auction problem. We denote the profit of a truthful mechanism $\mathcal{M}$ on input $\mathbf{b}$ by $p_{\mathcal{M}}(\mathbf{b})$. For randomized auction mechanisms $p_{\mathcal{M}}(\mathbf{b})$ is a random variable. A key question is how to evaluate the quality of the mechanism $\mathcal{M}$ as a profit maximizer. In Section 2 we present an upper bound on the profit achievable by any reasonable truthful mechanism given input bids $\mathbf{b}$. We use the bound to motivate a mechanism-independent quantity $p(\mathbf{b})$ against which we compare $p_{\mathcal{M}}(\mathbf{b})$. We then define the *competitive ratio* of $\mathcal{M}$ to be

$$\sup_{\mathbf{b}} \frac{p(\mathbf{b})}{\mathbf{E}[p_{\mathcal{M}}(\mathbf{b})]}.$$

[3]Bidders in different markets may also turn out to be indistinguishable from one another if the cost structure of providing services to them is identical.

Our goal is thus to construct truthful mechanisms that minimize this competitive ratio.

1.3 Cancellable Auctions

A natural approach to the design of generalized auction mechanisms that do not incur deficits is to run a basic auction on each market and then cancel the results if the revenue raised from the bidders does not exceed the costs that the auctioneer will have to pay to provide goods to the selected receivers. Consider, for example, the problem of selling a digital good via a truthful auction, where the cost of producing the good is C, and, once produced, the good can be duplicated at negligible cost (so that effectively the auctioneer has an unlimited supply of the good). In this case, the auctioneer (who will bear the burden of producing the good), does not even want to produce the good unless they can recover their cost of C.

We are thus motivated to introduce a stronger form of truthful auctions, cancellable auctions. Given a truthful auction $\mathcal{A}$ and a parameter C, we define an auction $\mathcal{A}_C$ as follows: On input $\mathbf{b}$, run $\mathcal{A}$ on $\mathbf{b}$. If the resulting revenue is at least C, return the outcome of $\mathcal{A}(\mathbf{b})$. Otherwise, *cancel* $\mathcal{A}$ by returning the outcome with no winners. We say that $\mathcal{A}$ is *cancellable* if, for any value of C, $\mathcal{A}_C$ is truthful. As we shall see in Section 3.1, not every truthful auction is cancellable.

1.4 Contributions

The contributions of this paper are twofold. We start by presenting a number of new results pertaining to the basic auction problem.

- We present a new, simple, and easy to analyze 4-competitive auction, and show that this auction is cancellable (See Section 3). Thus, as we shall see, it can be used as a building block in the design of mechanisms for generalized auction problems.

  The previous auctions known to be constant competitive for this problem (described in [7]) are more complicated to analyze and only large constant bounds (much greater than 4) are known for their competitive ratios. Also, they are not all cancellable.

- We prove a lower bound of 2 on the competitive ratio of any monotone truthful auction (For a definition of "monotone" and other pertinent details, see Section 3). We note that the class of monotone auctions is natural; all currently known competitive auctions are monotone.

In addition, we present a number of new results about generalized auction problems:

- We prove an upper bound on the profit achievable by any reasonable truthful mechanism, thereby motivating the performance measure we use for evaluating generalized auctions (See Section 2).

- We present a truthful mechanism that obtains a constant factor of the optimal truthful profit on *any* input set of bids for which (a) there is competition in each market, i.e., there is not one bidder whose bid dwarfs all other bids; and (b) there is a significant profit margin to be had, i.e., the optimal profit is at least a constant factor more than the cost incurred by the optimal allocation. Furthermore, on any set of bids not satisfying these conditions, the mechanism incurs no deficit (See Section 4).

  Thus, we obtain mechanisms that are profit maximizing for a broad class of generalized auction problems, including all the problems mentioned earlier in this section. For example, we obtain the first results on profit maximization for the conditional financing problem and for the multicast pricing problem in a worst-case, competitive analysis framework.[4] Although it is not the primary focus of our paper, the multicast pricing mechanism we present also has low network complexity, in the sense of Feigenbaum et al. [5].

1.5 Related Work

There has been a great deal of recent work at the intersection of game theory, economic theory and theoretical computer science [14, 17]. On the game theory and economics end, there is a large body of work on mechanism design (also known as implementation theory or theory of incentives) (see e.g., [12], chapter 23). One of the most important positive results in this field is the family of Vickrey-Clarke-Groves mechanisms [4, 8, 21].

Recent work in computer science pertaining to these fields has focused largely on merging the considerations of incentives with considerations of computational complexity (e.g., [9, 14, 15, 16]). One of the first examples of such work is the paper of Nisan and Ronen [15] in which the mechanism design framework is applied to some standard optimization problems in computer science, such as shortest paths and scheduling on unrelated machines.

Another line of research focuses on the complexity and design of cost sharing mechanisms. Feigenbaum, Papadimitriou and Shenker [5] study the distributed complexity of implementing two well-known cost sharing mechanisms, Marginal Cost and Shapley Value, for multicast transmission where routing is done over a subtree of a fixed universal tree. Jain and Vazirani [10] apply techniques from approximation algorithms to find cost sharing methods for approximate Steiner trees. The choice of mechanisms considered in [5, 10] is motivated by the desire to find a solution that is either budget-balanced or efficient, and hence there is no issue of maximizing the profit of the broadcaster.

Of course, auctions, be they traditional or combinatorial, have received a great deal of attention (see e.g., the surveys [11, 20]).

The work most closely related to our own is the work in [7], on which we build. Other work on profit maximization evaluated using worst-case analysis includes the work of Bar-Yossef et al. [3] in which an online version of the basic auction problem is studied in a competitive framework. Other interesting results in a similar framework include the work of Archer and Tardos [1], in which they present a mechanism which gets within a logarithmic factor of optimal for certain scheduling problems. Archer and Tardos [2] also show a $\Omega(n)$ lower bound on the competitive ratio for shortest path mechanisms even when there is close competition among the bidders.

[4]Recent work by Shenker and Vazirani studies profit maximization for the multicast pricing problem in a Bayesian setting [19]

2. PRELIMINARIES

In Section 1, we defined the generalized auction problem $\mathcal{A} = (\mathcal{S}, c(\cdot))$. A *mechanism* for this generalized auction problem is a mapping (perhaps randomized) from inputs to outputs. An input is a bid vector $\mathbf{b}$, where b_i, $1 \le i \le n$, is the bid input by the ith bidder. An output is an *allocation vector*, $\mathbf{x}$, and a *price vector*, $\mathbf{p}$. The allocation vector, $\mathbf{x}$, has $x_i = 1$ if bidder i receives the good and $x_i = 0$ otherwise and in the price vector, $\mathbf{p}$, p_i is the price bidder i pays. We will denote by $\mathbf{r}$ the allocation vector induced on the market segments, i.e., $r_j = 1$ if any bidders in S_j receive the good and $r_j = 0$ otherwise.

We assume that the ith bidder, $1 \le i \le n$, has a private *utility value* u_i and that each bidder's objective in choosing their bid value is to maximize their *profit*, defined as $u_i x_i - p_i$. We assume each bidder knows the mechanism, and that the bidders do not collude. We also impose a standard set of conditions on the auction mechanisms we consider:

- No Positive Transfers (NPT): $p_i \ge 0$. This precludes paying bidders to take an item.

- Voluntary Participation (VP): $p_i \le b_i$ if $x_i = 1$, and $p_i = 0$ if $x_i = 0$.

We say that a deterministic mechanism is *truthful* if each bidder, for any fixed choice of values for the other bids, can maximize their profit by bidding their utility value. We say that a randomized mechanism is truthful, if it is a probability distribution over deterministic mechanisms.[5]

Given a mechanism $\mathcal{M}$, for each input $\mathbf{b} \in \mathbb{R}^n$, the *profit to the auctioneer*, $p_{\mathcal{M}}(\mathbf{b})$, is the revenue from the bidders minus the cost of providing the goods, i.e.,

$$p_{\mathcal{M}}(\mathbf{b}) = \sum_i p_i - c(\mathbf{r}).$$

2.1 An Upper Bound on the Profit of Truthful Mechanisms

To motivate our performance measure, we will first prove a bound on the optimal auctioneer profit for any reasonable mechanism. To do this we need to first extend several notions from basic auctions to generalized auctions.

Bid-independence

We describe a useful characterization of truthful generalized auctions using 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*. Given a randomized function f on masked vectors (i.e., $f(\mathbf{b}_{-i})$ is a random variable), the *bid-independent generalized auction defined by* f on bid vector $\mathbf{b}$ is $\mathcal{M}_f$ (Auction 1).

Bid-independent auctions are truthful and vice versa.

THEOREM 2.1. *An auction is truthful if and only if it is the bid-independent auction defined by some function f.*

This is a straightforward generalization of the equivalent result for deterministic auctions in [7].

[5]This is a strong notion of truthfulness for randomized mechanisms. Weaker notions exist, but are not important for the results presented in this paper.

Auction 1 Bid-independent Auction ($\mathcal{M}_f$)

For each bidder i:

1. $v_i \leftarrow f(\mathbf{b}_{-i})$.
2. If $v_i \le b_i$, set $x_i \leftarrow 1$ and $p_i \leftarrow v_i$.[6] (Bidder i wins.)
3. Otherwise, set $x_i = p_i = 0$. (In this case, we say that bidder i is rejected.)

Note that $f(\mathbf{b}_{-i})$ and $f(\mathbf{b}_{-j})$ need not be independent.

Monotonicity

Using standard terminology, we say that random variable X *dominates* random variable Y if for all x

$$\mathbf{Pr}[X \ge x] \ge \mathbf{Pr}[Y \ge x]\,.$$

We are now ready to define monotone generalized auctions.

DEFINITION 2.2. *A generalized auction, $\mathcal{M}$, is monotone on markets $\mathcal{S}$ if it is defined by a bid-independent function f with the property that for any bid vector $\mathbf{b}$, and any pair of bidders i and j within the same market such that $b_i \le b_j$ the random variable $f(\mathbf{b}_{-i})$ dominates the random variable $f(\mathbf{b}_{-j})$.*

To get a feel for this definition, observe that when $b_i \ge b_j$ the bids visible in the masked vector $\mathbf{b}_{-j}$ are the same as those visible in the masked vector $\mathbf{b}_{-i}$ except for the fact that the larger bid b_i is visible in $\mathbf{b}_{-j}$ whereas the smaller bid b_j is visible in $\mathbf{b}_{-i}$. Intuitively, monotonicity means that the bid-independent function, upon seeing a set of higher bid values outputs higher threshold prices.

Optimal fixed-pricing

We first recall the definition of optimal fixed pricing for basic auctions [7]:

DEFINITION 2.3. *Let $\mathbf{b}$ be a bid vector and let $\overline{b}_i$ be the ith largest component of the vector $\mathbf{b}$. Then the* optimal fixed price revenue *on $\mathbf{b}$, $\mathcal{F}(\mathbf{b})$, is*

$$\mathcal{F}(\mathbf{b}) = \max_i i\overline{b}_i.$$

Note that $\mathcal{F}(\mathbf{b})$ can be $\Theta(\frac{1}{\log n}\sum_i b_i)$, and therefore a $\log n$ factor from the sum of the bids, an obvious upper bound on the revenue for a multi-priced auction.[7] $\mathcal{F}(\mathbf{b})$ is the optimal profit for any mechanism that uses a single selling price for all winners.

We now extend this to general auctions:

DEFINITION 2.4. *The optimal profit for any selling mechanism for $\mathcal{A} = (\mathcal{S}, c(\cdot))$ that uses a single price for each market is*

$$\mathcal{F}_{\mathcal{A}}(\mathbf{b}) = \max_{\mathbf{r} \in \{0,1\}^m} \left(\sum_{1 \le j \le m} r_j \mathcal{F}(\mathbf{b}_{S_j}) - c(\mathbf{r}) \right).$$

where $\mathbf{b}_{S_j}$ is the restriction of $\mathbf{b}$ to the jth market S_j.

[7]Consider, for example, n bids with $b_i = 1/i$.

Profit upper bound

THEOREM 2.5. *Let $\mathcal{M}$ be any monotone auction for the basic auction problem on bids* $\mathbf{b}$*, and let* $\mathcal{W}$ *be the event that there is at least one winner, i.e., there exists* i *such that* $x_i = 1$*. The revenue* $\mathcal{R} = \sum_i p_i$ *of* $\mathcal{M}$ *on input* $\mathbf{b}$ *satisfies:*

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

PROOF. Let f be the bid-independent function defining $\mathcal{M}$. Let q be the probability that the event $\mathcal{W}$ occurs. We define $g_i(x)$ as follows:

$$\begin{aligned} g_i(x) &= \mathbf{Pr}[f(\mathbf{b}_{-i}) \leq x \mid \mathcal{W}] \\ &= \tfrac{1}{q}\mathbf{Pr}[f(\mathbf{b}_{-i}) \leq x \cap \mathcal{W}] \end{aligned}$$

For $x \leq b_i$ bidder i is a winner and thus $f(\mathbf{b}_{-i}) \leq x$ implies event $\mathcal{W}$. So we can conclude that for $x \leq b_i$,

$$g_i(x) = \tfrac{1}{q}\mathbf{Pr}[f(\mathbf{b}_{-i}) \leq x].$$

By the monotonicity of f, for all i and j, g_i satisfies the property that if $b_i \geq b_j$ then $g_i(x) \geq g_j(x)$ for $x \leq b_j$.

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 using $g_i^{-1}(U)$ 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 $\mathcal{R}_U$ 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 conditioned on $\mathcal{W}$:

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

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

$$\mathbf{E}[\mathcal{R}_U] = \mathbf{E}[\mathcal{R}|\mathcal{W}]\,.$$

We complete the proof by showing that the expected revenue from our thought experiment $\mathbf{E}[\mathcal{R}_U]$ 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 $g_k(x) \leq g_j(x)$ for all j such that $b_j \geq b_k$, and $g_k(x)$ and $g_j(x)$ are monotone non-decreasing functions, we must have $g_j^{-1}(u) \leq g_k^{-1}(u) \leq b_k \leq b_j$ and therefore all bidders with bid values at least b_k win at a price at most b_k. Thus, the revenue, $\mathcal{R}_u$, 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 $\mathbf{E}[\mathcal{R}_U] \leq \mathcal{F}(\mathbf{b})$. ∎

COROLLARY 2.6. *The expected profit for any monotone basic auction on any set of bids* $\mathbf{b}$ *is at most* $\mathcal{F}(\mathbf{b})$.[8]

We now give an upper bound on profit for generalized auctions.

THEOREM 2.7. *Let* $\mathcal{A} = (\mathcal{S}, c(\cdot))$ *be a generalized auction problem. For any truthful mechanism* $\mathcal{M}$ *that is monotone on* $\mathcal{S}$*, and for any input bid vector* $\mathbf{b}$*, we give an upper bound* $p_{\mathcal{M}}(\mathbf{b})$*, the profit of* $\mathcal{M}$ *on bid set* $\mathbf{b}$ *as:*

$$\mathbf{E}[p_{\mathcal{M}}(\mathbf{b})] \leq \mathcal{F}_{\mathcal{A}}(\mathbf{b}).$$

[8] A similar result was claimed in [7]; however, the proof was incorrect.

PROOF. Define the following:

- $q_{\mathbf{r}}$ with $\mathbf{r} = (r_1, \ldots, r_m) \in \{0,1\}^m$ – the probability that the result of the auction is the allocation $\mathbf{r}$ to the markets.

For $1 \leq j \leq m$ define:

- q_j – the probability that there is at least one winner in market S_j. Note that $q_j = \sum_{\mathbf{r}\in\{0,1\}^m} r_j q_{\mathbf{r}}$.
- $\mathcal{F}_j = \mathcal{F}(\mathbf{b}_{S_j})$ – the optimal fixed price revenue for market j.
- $\mathcal{R}_j$ – the revenue collected from the jth market.

Note that from Theorem 2.5 we have

$$\begin{aligned} \mathbf{E}[\mathcal{R}_j] &= q_j \mathbf{E}[\mathcal{R}_j \mid r_j = 1] \\ &\leq q_j \mathcal{F}_j \end{aligned} \tag{1}$$

Now we bound the expected profit:

$$\begin{aligned} \mathbf{E}[p_{\mathcal{M}}(\mathbf{b})] &= \mathbf{E}\Big[\sum\nolimits_j \mathcal{R}_j - c(\mathbf{r})\Big] \\ &= \sum\nolimits_j \mathbf{E}[\mathcal{R}_j] - \mathbf{E}[c(\mathbf{r})]\,. \end{aligned}$$

By equation (1),

$$\begin{aligned} \mathbf{E}[p_{\mathcal{M}}(\mathbf{b})] &\leq \sum\nolimits_j q_j \mathcal{F}_j - \mathbf{E}[c(\mathbf{r})] \\ &= \sum\nolimits_j \mathcal{F}_j \left(\sum\nolimits_{\mathbf{r}} r_j q_{\mathbf{r}}\right) - \sum\nolimits_{\mathbf{r}} q_{\mathbf{r}} c(\mathbf{r}) \\ &= \sum\nolimits_{\mathbf{r}} q_{\mathbf{r}} \left(\sum\nolimits_j r_j \mathcal{F}_j - c(\mathbf{r})\right). \end{aligned}$$

This is a convex combination, so we have

$$\begin{aligned} \mathbf{E}[p_{\mathcal{M}}(\mathbf{b})] &\leq \max_{\mathbf{r}\in\{0,1\}^m} \left(\sum_{1\leq j\leq m} r_j \mathcal{F}_j - q_{\mathbf{r}} c(\mathbf{r})\right) \\ &= \mathcal{F}_{\mathcal{A}}(\mathbf{b}). \end{aligned}$$

∎

The proof just presented may seem at first glance to be more complicated than necessary. However, we note that the most obvious approach of using the equation

$$\mathbf{E}[p_{\mathcal{M}}(\mathbf{b})] = \sum_{\mathbf{r}} q_{\mathbf{r}} \left(\sum_j \mathbf{E}[\mathcal{R}_j|\mathbf{r}] - c(\mathbf{r})\right)$$

and then showing that $\mathbf{E}[\mathcal{R}_j|\mathbf{r}] \leq \mathcal{F}_j$ fails: $\mathbf{E}[\mathcal{R}_j|\mathbf{r}]$ can in fact significantly exceed $\mathcal{F}_j$.

2.2 Competitive Mechanisms

The fact that $\mathcal{F}_{\mathcal{A}}(\mathbf{b})$ is an upper bound on the profit of any monotone truthful auction on input set $\mathbf{b}$ motivates the evaluation of truthful mechanisms by considering their competitive ratio relative to $\mathcal{F}_{\mathcal{A}}$, i.e., the supremum over all bid vectors of $\mathcal{F}_{\mathcal{A}}(\mathbf{b})/p_{\mathcal{M}}(\mathbf{b})$. Unfortunately, in many cases, there is no constant bound achievable on this competitive ratio. For example, for the basic unlimited supply auction problem with no costs, $\mathcal{F}_{\mathcal{A}}(\mathbf{b})$ is equal to the optimal fixed price revenue $\mathcal{F}(\mathbf{b})$. However, it is easy to show that no randomized unlimited supply auction can obtain a constant fraction of $\mathcal{F}(\mathbf{b})$. Intuitively, if there is one very high bidder that completely dominates all other bidders, there is no way to truthfully extract a constant fraction of his bid value [6]. In contrast, as we show below, it is possible to obtain a constant fraction of $\mathcal{F}^{(2)}(\mathbf{b})$, where

$$\mathcal{F}^{(2)}(\mathbf{b}) = \max_{i \geq 2} i\overline{b}_i,$$

i.e., the optimal fixed price revenue assuming at least two items are sold (again, $\overline{b}_i$ is the ith largest bid).

As the basic auction is a special case of the general auction, it is not possible to be constant competitive with $\mathcal{F}_{\mathcal{A}}(\mathbf{b})$. Furthermore, we conjecture that for the fixed cost unlimited supply auction where the cost of producing the good is C and $\mathcal{F}_{\mathcal{A}}(\mathbf{b}) = \max(0, \mathcal{F}(\mathbf{b}) - C)$ it is not possible to even attain a constant fraction of $\max(0, \mathcal{F}^{(2)}(\mathbf{b}) - C)$ as this difference gets arbitrarily small. Intuitively, it is much easier to be competitive with gross profit than it is to be competitive with the net profit.

We are thus motivated to define the following weaker notion of competitiveness:

Definition 2.8. *For generalized auction problems, we say that a truthful mechanism $\mathcal{M}$ is β-*competitive *if, for all bid vectors* $\mathbf{b}$*,*

$$\mathbf{E}[p_{\mathcal{M}}(\mathbf{b})] \geq \frac{\text{profit}_\beta(\mathbf{b})}{\beta},$$

where

$$\text{profit}_\beta(\mathbf{b}) = \max_{\mathbf{r} \in \{0,1\}^m} \left(\sum_{1 \leq i \leq m} r_i \mathcal{F}^{(2)}(\mathbf{b}_{S_i}) - \beta c(\mathbf{r}) \right),$$

As noted previously, we cannot be constant competitive against $\mathcal{F}_{\mathcal{A}}(\mathbf{b})$. Nonetheless, for a large class of inputs, achieving a constant fraction of $\mathcal{F}_{\mathcal{A}}(\mathbf{b})$ implies achieving a constant fraction of $\text{profit}_\beta(\mathbf{b})$ and vice versa. This is the class of bids for which

- For each market S_j, $\mathcal{F}(\mathbf{b}_{S_j}) = \mathcal{F}^{(2)}(\mathbf{b}_{S_j})$, meaning there is no single bidder whose bid completely dominates all others in that market, and
- The optimal profit margin is a constant factor of the cost of the optimal allocation.

Thus, for an interesting class of problems and inputs to these problems, we will present mechanisms that obtain profit that is within a constant factor of optimal. An interesting problem left open by our work is whether or not there is a stronger performance measure one can compare against and still obtain strong competitiveness guarantees.

2.3 Randomized Mechanisms

In [7] it is proven that no deterministic symmetric[9] mechanism for the basic auction problem is constant competitive. Thus, we are compelled to use randomized mechanisms, and shall do so for the remainder of the paper.

[9] An auction is symmetric if its outcome is the same for any permutation of the input bids.

3. BASIC AUCTION RESULTS

In this section, we consider the basic auction problem assuming the auctioneer has an unlimited supply of items.[10] All the results presented in this section can be generalized easily to hold for the case where the auctioneer only has a fixed number, k, of items. For a k-item auction, simply reject all but the k highest bidders and run an unlimited supply auction on the remaining bids.

3.1 Cancellability

Cancellable auctions are crucial building blocks in the design of mechanisms for more general auctions. As discussed in the introduction, we define cancellable auctions as follows.

Definition 3.1. *Given a truthful auction $\mathcal{M}$ and a parameter C, we define the auction $\mathcal{M}_C$ as follows: On input* $\mathbf{b}$*, run $\mathcal{M}$ on* $\mathbf{b}$*. If the resulting revenue is at least C, return the outcome of $\mathcal{M}(\mathbf{b})$. Otherwise,* cancel *$\mathcal{M}$ by returning the outcome with no winners. We say that $\mathcal{M}$ is* cancellable *if and only if, for any value of C, $\mathcal{M}_C$ is truthful.*

There are two natural ways to prove that an auction is cancellable: The first is to show that for each C, there is a bid-independent function f such that the outcome (as a random variable) of the bid-independent auction defined by f is the same as that of $\mathcal{M}_C$. A second approach is to show directly that in $\mathcal{M}_C$ any bidder's profit is maximized by bidding their utility value. Specifically it suffices to use the following:

Proposition 3.2. *The truthful auction $\mathcal{M}$ is cancellable if for any bid vector $\mathbf{b}_{-i}$ and any bidder i with utility u_i, if bidder i wins upon bidding u_i, then $R_{u_i} \geq R_b$ for any b such that $x_i = 1$ when bidder i bids b. For the fixed values of $\mathbf{b}_{-i}$, R_b denotes the revenue of $\mathcal{M}$ when bidder i bids b.*

To see this, observe that, by assumption, $\mathcal{M}$ is truthful, and thus, bidder i's profit before the possible cancellation is maximized by bidding their utility value. To ensure that bidder i's profit is still maximized after cancellation by bidding u_i, it must be that if the auction is cancelled when bidder i bids u_i, then bidder i loses when bidding any value that results in the auction not being cancelled. Another way to put this is that a truthful auction is cancellable if its revenue is not a function of the bid values of the winning bids.

We next observe that the notion of a cancellable auction is strictly stronger than that of a truthful auction.

Observation 3.3. *Not all truthful auction mechanisms are cancellable.*

In particular, one variant of the class of constant-competitive auctions presented in [7], the *Dual-price Sampling Optimal Threshold* auction (Auction 2, below), is not cancellable.[11] DSOT is obviously truthful, as it is bid-independent. We observe now that the auction obtained by running DSOT and cancelling the outcome if the revenue is less than C is not truthful.

[10] In terms of the generalized auction problem, this means that there is only one market containing all the bidders, and the cost function is identically 0.

[11] In fact, none of the auctions developed in [7] are cancellable except for the *Random Sampling Optimal Threshold* auction (RSOT).

Consider $\mathbf{b} = \{1, 1, 1, \ldots, 1, h\}$. Clearly, for h sufficiently large, h is the optimal threshold for its partition: bids in the other partition will be rejected and the revenue of the auction is close to $n/2$. However, if the high bidder bids 1 instead of h, $\mathbf{b}^* = \{1, 1, 1, \ldots, 1, 1\}$ and all bidders get $x_i = 1$ and $p_i = 1$. Thus the auction revenue is n.

Suppose the auction is to be cancelled if its revenue is less than $C = \frac{3}{4}n$. Then,

- the high bidder's profit is 0 when bidding h, and
- the high bidder's profit is $h - 1$ when bidding 1.

Thus DSOT is not cancellable.

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

For any input $\mathbf{b}$:

1. Partition $\mathbf{b}$ into two sets, $\mathbf{b}_{S'}$ and $\mathbf{b}_{S''}$, by flipping a fair coin for each bid.
2. Let p' (resp. p'') be the optimal fixed price for the bids in $\mathbf{b}_{S'}$ (resp. $\mathbf{b}_{S''}$).
3. Use p' as a threshold for bids in S'' and p'' for bids in S' (I.e., for $i \in S'$ if $b_i \geq p''$ then set $x_i \leftarrow 1$ and $p_i \leftarrow p'$; otherwise set $x_i \leftarrow 0$ and $p_i \leftarrow 0$).

3.2 The Sampling Cost Sharing Auction

The main result of this section is a competitive truthful auction that is cancellable. It it simple, easy to analyze, and achieves a competitive ratio that is better than that known for any other auction.

We first review a standard cost sharing mechanism [13, 18]. Given bids $\mathbf{b}$ and cost C, this mechanism finds a subset of the bidders to share the cost C, if possible. 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 can equally share the cost C. Charge each C/k.

Important properties of this mechanism are as follows.

- If $C \leq \mathcal{F}(\mathbf{b})$, CostShare$_C$ has revenue C; otherwise it has no revenue and the outcome is $\mathbf{x} = \mathbf{0}$ and $\mathbf{p} = \mathbf{0}$.
- CostShare$_C$ is truthful.[12]

A formal description of the *Sampling Cost Sharing* auction is given in Auction 3. As we will show, this auction has the property that it is 4-competitive and cancellable. SCS computes the optimal fixed price revenues for two random partitions of the bidders and has each partition cost share the other partitions optimal revenue. Typically, these optimal fixed price revenues will be different and the partition with the lesser revenue will be completely rejected as it cannot afford to cost share the optimal revenue of partition with the higher optimal revenue. In the case where the optimal revenues are the same, we could either allow both partitions to have winners or impose some tie breaking criterion to make sure that only one of the partitions has winners. If we allow both partitions to have winners, the resulting mechanism, while being truthful and 4-competitive, is not cancellable.[13] Thus, if we want a cancellable auction, we must do some sort of tie breaking. The technique we choose, as described in the formal definition of SCS, is to use the total order, $\prec$, on fixed price revenues. We omit a more detailed explanation of this tie breaking technique from this extended abstract.

Auction 3 Sampling Cost Sharing Auction (SCS)

1. Partition bids $\mathbf{b}$, into two sets S' and S'', by flipping a fair coin for each bid. Let the resulting bid vectors be $\mathbf{b}_{S'}$ and $\mathbf{b}_{S''}$.
2. Compute $\mathcal{F}' = \mathcal{F}(\mathbf{b}_{S'})$ and $\mathcal{F}'' = \mathcal{F}(\mathbf{b}_{S''})$, the optimal fixed price revenues for $\mathbf{b}_{S'}$ and $\mathbf{b}_{S''}$, respectively.
3. Compute the auction results by running CostShare$_{\mathcal{F}''}$ on $\mathbf{b}_{S'}$ and CostShare$_{\mathcal{F}'}$ on $\mathbf{b}_{S''}$.

We impose a total ordering, "$\prec$", on values of the form "kb_i" that respects their natural partial ordering given by "$<$". We define $\prec$ as:

$$kb_i \prec \ell b_j \iff kb_i < \ell b_j \vee kb_i = \ell b_j \wedge i < j.$$

Note that $\mathcal{F}'$ and $\mathcal{F}''$ are of the form "kb_i" for some k and i, so using this total ordering in Steps 2 and 3 guarantees $\mathcal{F}' \neq \mathcal{F}''$.

THEOREM 3.4. *SCS is 4-competitive, and this bound is tight.*

PROOF. We begin by observing that the auction revenue is $\mathcal{R} = \min(\mathcal{F}', \mathcal{F}'')$. Suppose, without loss of generality, that $\mathcal{F}' < \mathcal{F}''$. Then CostShare$_{\mathcal{F}''}$ on $\mathbf{b}_{S'}$ will reject all bids in $\mathbf{b}_{S'}$. However, CostShare$_{\mathcal{F}'}$ on $\mathbf{b}_{S''}$ will be able to achieve revenue $\mathcal{F}'$.

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

[12] For example, we can exhibit a function f such that CostShare$_C$ is equivalent to the bid-independent auction defined by f:

1. Let $\mathbf{b}'$ be $\mathbf{b}_{-i}$ with b'_i set to value C.
2. Find the largest k such that the highest k bidders can equally share the cost C.
3. Return C/k as the threshold for bidder i.

It is easy to verify that the auction defined by f is in fact CostShare$_C$.

[13] For example, consider two bidders with utilities $u_1 = 4$ and $u_2 = 2$ and a cost $C = 3$. Assume we sample with $S' = \{1\}$ and $S'' = \{2\}$. If both bidders bid their utility we have $\mathcal{F}' = 4$ and $\mathcal{F}'' = 2$ and the revenue is $\mathcal{R} = 2$. Since the revenue is less than the cost C, the auction is cancelled and the first bidder's profit is zero. If we allow ties, the first bidder could lower their bid to $b_1 = 2$ causing $\mathcal{F}' = \mathcal{F}'' = 2$. This results in both bidders winning and yields a revenue of $\mathcal{R} = 4$. In this case, the auction is not cancelled and the first bidder's profit is two.

$\mathcal{F}(\mathbf{b}_{S''}) \geq pk''$. Therefore,

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

Thus, the competitive ratio

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

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

To see that the bound 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 negligibly small. In this case $\mathcal{F} = \mathcal{F}^{(2)} = 2h$, whereas the expected revenue of the SCS auction is $h \cdot \mathbf{Pr}[\text{two high bids are split}] = h/2 = \mathcal{F}/4$. ■

Finally, we prove the following.

LEMMA 3.5. *SCS is a cancellable auction.*

PROOF. Let SCS_C be the auction that runs SCS and cancels the outcome if its revenue is not at least C. By definition, SCS is cancellable if and only if SCS_C is truthful for all C. We prove this using Proposition 3.2.

Consider any fixed outcome of the coin flips of SCS, and any bid vector $\mathbf{b}_{-i}$. In addition, suppose that bidder i's utility value is u_i. For this particular execution, suppose that bidder i is in S'. If bidder i wins in this execution of SCS (prior to the possible cancellation) then $\mathcal{F}'' \prec \mathcal{F}'$. Now suppose that bidder i changes his bid, resulting, possibly, in a new value of $\mathcal{F}'$. If we still have $\mathcal{F}'' \prec \mathcal{F}'$, then the total revenue of the auction is unchanged, equal to $\mathcal{F}''$. Otherwise, if we now have $\mathcal{F}' \prec \mathcal{F}''$, then bidder i loses. Thus, in terms of Proposition 3.2, $R_{u_i} \geq R_b$ for any b such that $x_i = 1$ when bidder i bids b. ■

3.3 Lower bound on Competitive Ratio

We say that auction defined by bid-independent funciton f is *scale-invariant* if, for all i and all x, $\mathbf{Pr}[f(\mathbf{b}_{-i}) \geq x] = \mathbf{Pr}[f(c\mathbf{b}_{-i}) \geq cx]$. In this section we show that any scale-invariant auction has a competitive ratio of at least 2. When bids are restricted to be within a certain range (e.g., in the range $[1, M]$) or are restricted to be powers of α for some $\alpha > 1$, it is possible to design non-scale-invariant auctions that do better. However, for arbitrary nonnegative bids, we conjecture that for any auction with competitive ratio c, there is a scale-invariant auction with competitive ratio c. If this conjecture is true, then the competitive ratio of SCS is within a factor of 2 of optimal.

We say that an auction $\mathcal{M}$ is *symmetric* if $\mathcal{M}(\mathbf{b}) = \mathcal{M}(\pi(\mathbf{b}))$ for any permutation $\pi(\mathbf{b})$ of the bid vector $\mathbf{b}$.

LEMMA 3.6. *Let $\mathcal{M}$ be an auction with competitive ratio c. Then there is a randomized symmetric auction $\mathcal{M}'$ with competitive ratio c.*

PROOF. Given an auction $\mathcal{M}$, we define $\mathcal{M}'$ as follows: Apply a random permutation π to the input bid vector $\mathbf{b}$ and run $\mathcal{M}$ on $\pi(\mathbf{b})$. Clearly $\mathcal{M}$ is symmetric. Furthermore, applying a permutation to the input vector does not change the value of $\mathcal{F}^{(2)}$. By the choice of $\mathcal{M}$, for any $\mathbf{b}$ and π, the (expected) revenue of $\mathcal{M}$ on $\pi(\mathbf{b})$ is at least $\mathcal{F}^{(2)}/c$ and therefore the expected revenue of $\mathcal{M}'$ on $\mathbf{b}$ is at least $\mathcal{F}^{(2)}/c$. ■

THEOREM 3.7. *The competitive ratio of any randomized scale-invariant auction is at least two.*

PROOF. We will show this result for the special case where there are only two bidders. Recall that any truthful auction is a bid-independent auction $\mathcal{M}_f$ for some f, where f is randomized function on masked vectors $\mathbf{b}_{-i}$. Lemma 3.6 allows us to consider, without loss of generality, symmetric auctions only. For symmetric two bidder auctions, given bids b_1 and b_2, $f(b_2)$ determines the price for b_1 and vice versa. For scale-invariant auctions, $f(b_2)$ is completely determined by $f(1)$ scaled appropriately by b_2.

Note that $\mathcal{F}^{(2)}(\mathbf{b}) = 2\min(b_1, b_2)$ for an auction with only two bidders. Given bids of the form $\mathbf{b} = (1, h)$ for any $h \geq 1$, $\mathcal{F}^{(2)}(\mathbf{b}) = 2$. Consider the deterministic bid-independent and scale-invariant function $g(x) = x$. For two bidders, the corresponding auction $\mathcal{M}_g$ is trivially 2-competitive as its revenue is at least $\min(b_1, b_2)$. We show that this auction achieves the best (smallest) competitive ratio among scale-invariant auctions.

We focus here on the case that $f(1)$ is a finite, discrete random variable. The general case follows by taking the limit as the number of different values taken on by the random variable goes to infinity and by approximating any continuous distribution with a discrete one. We will now show how from any finite, discrete random variable $f(1)$ representing a two bidder auction, we can inductively construct a sequence of auctions with non-increasing competitive ratio ultimately arriving at $\mathcal{M}_g$. This would imply that $\mathcal{M}_f$'s competitive ratio is no better that of $\mathcal{M}_g$ which is 2.

Let $p(x) = \mathbf{Pr}[f(1) = x]$ and let $\mathcal{R}_f(h)$ denote the expected revenue for $\mathcal{M}_f$ on bids $\{1, h\}$. That is,

$$\mathcal{R}_f(h) = \underbrace{\sum_{0 \leq x \leq 1} xp(x/h)}_{\text{from } 1} + \underbrace{\sum_{0 \leq x \leq h} xp(x)}_{\text{from } h}$$

The proof of the theorem is by induction on m, where m is the number of values x such that $x \neq 1$ and $p(x) > 0$. For the base case, $m = 0$, and then, trivially $f = g$. Given nonzero m, we can construct a new auction $\mathcal{M}_{f''}$ with $p''(x) = \mathbf{Pr}[f''(1) = x]$ that has less than m values x such that $x \neq 1$ and $p''(x) > 0$. We do this in two steps. Let $\underline{x}$ denote the lesser of the value one and the smallest value of x such that $p(x)$ is non-zero. Likewise, let $\bar{x}$ denote the larger of the value one and the largest value of x such that $p(x)$ is non-zero.

The basic inductive step of the proof will consist of first constructing p' from p with the property that $1/\underline{x}' = \bar{x}'$. We then move probability mass from $\underline{x}'$ and $\bar{x}'$ to 1 to get p'' and the auction $\mathcal{M}_{f''}$ with $m'' < m$. To show that the competitive ratio of $\mathcal{M}_f$ is at least that of $\mathcal{M}_{f'}$ it suffices to show that any value of h', we can find a value h such that $\mathcal{R}_{f'}(h') \geq \mathcal{R}_f(h)$.

We begin by constructing p'. There are three cases.

Case 1, $1/\underline{x} = \bar{x}$: Define p' as p.

$$p'(x) = p(x)$$

Case 2, $1/\underline{x} < \bar{x}$: For example,

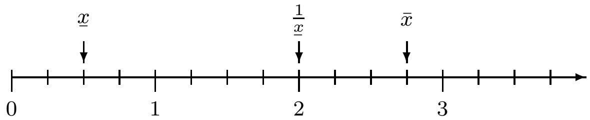

Define p' as p but with all weight above $1/\underline{x}$ moved down to be at $1/\underline{x}$. That is,

$$p'(x) = \begin{cases} 0 & \text{if } x > 1/\underline{x} \\ \sum_{x' \geq 1/\underline{x}} p(x') & \text{if } x = 1/\underline{x} \\ p(x) & \text{otherwise.} \end{cases}$$

To show that $\mathcal{M}_{f'}$ defined by this p' has a competitive ratio that is at most that of $\mathcal{M}_f$, we consider all $h' \geq 1$ and show that there exists an h such that $\mathcal{R}_{f'}(h') \geq \mathcal{R}_f(h)$.

- For $h' < 1/\underline{x}$, $\mathcal{R}_{f'}(h') = \mathcal{R}_f(h')$ because $p(x) = p'(x)$ for $x \in [0, 1/\underline{x})$. Thus, choose $h = h'$.
- For $h' = 1/\underline{x}$, $\mathcal{R}_{f'}(h') = \mathcal{R}_f(h') + h(p'(h) - p(h)) > \mathcal{R}_f(h')$. Thus, choose $h = h'$.
- If $h' > 1/\underline{x}$ then bid 1 is rejected by both $\mathcal{M}_{f'}$ and $\mathcal{M}_f$ (since for all $y \leq 1$, $p(y/h') = p'(y/h') = 0$). Note that since $p(x)$ is non-zero on finite number of x there exists an ϵ such that with $h = 1/\underline{x} + \epsilon$,

$$\mathcal{R}_f(h) = \sum_{0 \leq x}^{1/\underline{x}} x p(x) \leq \mathcal{R}_{f'}(h')$$

For all $h' > 1/\underline{x}$ choose this h.

Case 3, $1/\underline{x} > \bar{x}$: Define p' as p but with all weight below $1/\bar{x}$ moved up to be at $1/\bar{x}$. That is,

$$p'(x) = \begin{cases} 0 & \text{if } x < 1/\bar{x} \\ \sum_{x' \leq 1/\bar{x}} p(x') & \text{if } x = 1/\bar{x} \\ p(x) & \text{otherwise.} \end{cases}$$

Showing that the competitive ratio of $\mathcal{M}_{f'}$ is at most that of $\mathcal{M}_f$ is similar in spirit to Case 2 and we omit the details.

After any of case 1, 2, or 3 above, we are left with p' such that $1/\underline{x} = \bar{x}$ and the corresponding auction $\mathcal{M}_{f'}$ has at most the same competitive ratio as $\mathcal{M}_f$. Now we seek to reduce the number of values x such that $x \neq 1$ and $p(x) > 0$.

Define p'' as p' but with as much as possible of the mass from $\underline{x}$ and $\bar{x}$ moved to be at 1, balanced so as to keep the expected value from changing. Thus, we move mass q from $\underline{x}$ and mass $q/\bar{x}$ from $\bar{x}$ with $q = \min(p'(\underline{x}), p'(\bar{x})\bar{x})$ so that either $\underline{x}$ or $\bar{x}$ will have mass zero in p'', thus reducing the number of values x such that $x \neq 1$ and $p'(x) > 0$.

$$p''(x) = \begin{cases} p'(x) - q & \text{if } x = \underline{x} \\ p'(x) - q/\bar{x} & \text{if } x = \bar{x} \\ p'(x) + q + q/\bar{x} & \text{if } x = 1 \\ p'(x) & \text{otherwise.} \end{cases}$$

Showing that the competitive ratio of $\mathcal{M}_{f''}$ is at most that of $\mathcal{M}_{f'}$ is also similar in spirit to Case 2 and we again omit the details. ■

4. APPLICATIONS OF CANCELLABLE AUCTIONS

In this section we show how cancellable auctions can be used as a building block for the construction of profit optimizing mechanisms. For generalized auction problems, we define the following *Local Sampling Cost Sharing* auction (Auction 4).

Auction 4 Local Sampling Cost Sharing auction (LSCS)

On input $\mathbf{b}$:

1. Run SCS on each market, $\mathbf{b}_{S_j}$, to get revenue $\mathcal{R}_j$.
2. Compute the maximizing (or approximate) market allocation $\mathbf{r}^* = \operatorname{argmax}_{\mathbf{r}} \sum_j r_j \mathcal{R}_j - c(\mathbf{r})$.
3. For each j, if $r_j^* = 0$, cancel auction on S_j. Else collect $\mathcal{R}_j$ in revenue from market S_j.

LEMMA 4.1. *The Local Sampling Cost Sharing auction is truthful.*[14]

This proof is immediate from the cancellability of SCS.

THEOREM 4.2. *LSCS is 4-competitive.*

PROOF. Since SCS is 4-competitive, for all j we can expect at least $\mathcal{F}^{(2)}(\mathbf{b}_{S_j})/4$ from the jth market. Consider the allocation $\mathbf{r}'$ used by $\text{profit}_4(\mathbf{b})$, that is, $\mathbf{r}'$ maximizes $\max_{\mathbf{r}} \sum_j r_j \mathcal{F}(\mathbf{b}_{S_j}) - 4c(\mathbf{r})$.

Our revenue is

$$p_{\text{LSCS}}(\mathbf{b}) = \max_{\mathbf{r}} \sum_j \mathcal{R}_j r_i - c(\mathbf{r}) \geq \sum_j \mathcal{R}_j r'_i - c(\mathbf{r}')$$

and

$$\mathbf{E}\left[\sum_j \mathcal{R}_j r'_i - c(\mathbf{r}')\right] \geq \sum_j \mathcal{F}^{(2)}(\mathbf{b}_{S_j}) r'_i/4 - c(\mathbf{r}') = \text{profit}_4(\mathbf{b})/4$$

and so

$$\mathbf{E}[p_{\text{LSCS}}(\mathbf{b})] \geq \text{profit}_4(\mathbf{b})/4.$$

■

COROLLARY 4.3. *LSCS is 4-competitive for the multicast problem.*

We note that LSCS for the multicast problem of [5] can be implemented with low network complexity, 2 messages per link, using the natural prize-collection algorithm for trees (E.g. [5]).

COROLLARY 4.4. *Consider the algorithm LSCS′ that differs from LSCS in that rather than actually compute the optimal allocation, it makes use of an α approximation to the optimal allocation. LSCS′ is 4α-competitive.*

[14] In fact, any cancellable auction could be used in place of SCS and the resulting local mechanism would be truthful.

As we have seen, the LSCS auction works very well for generalized auction problems when each market has size at least 2 and there is competition in each market.

These assumptions might be valid for multicast in the Internet today where the content provider is charged for usage of the backbone, while consumers are located at single ISPs and are already paying a flat rate for their service. In this case, the consumers at each ISP would form markets and there would possibly be a large number of them at each.

A key open question is how well a more global auction mechanism can do when these assumptions do not hold. For example, there is a relatively simple global mechanism for the basic market segmentation problem (where there are m markets but the cost, $c(\mathbf{r})$, is zero for all $\mathbf{r}$). It remains open whether this or any other more global mechanism works well in any quantifiable sense.

5. CONCLUSIONS

For the basic auction problem we have presented a number of results. The most important of these are that the Sampling Cost Sharing auction is 4-competitive and cancellable and that no scale-invariant auction can achieve a competitive ratio less than two. A natural resulting open question is that of determining the optimal competitive ratio for any basic auction. However, even for the two bidder case, a "proof from the book" that no auction (scale-invariant or not) can achieve a competitive ratio of better than two would be an interesting result.

For the fixed cost unlimited supply auction, i.e., one market with cost C if there are any winners, we have shown that it is possible to get a constant fraction of $\mathcal{F}^{(2)}(\mathbf{b}) - 4C$ in worst case. We conjecture that this scaling of C is necessary in the sense that it is not possible to get a constant fraction of $\mathcal{F}^{(2)}(\mathbf{b}) - C$ in worst case. This result seems likely because $\mathcal{F}^{(2)}(\mathbf{b}) - C$ can be arbitrarily small compared to $\mathcal{F}^{(2)}(\mathbf{b})$ and we know it is not possible to be better than 2-competitive against $\mathcal{F}^{(2)}(\mathbf{b})$.

For general auctions we gave a local mechanism that is competitive if there is competition amongst the bidders in each market and if there is a significant profit margin. This solution can be applied directly to the multicast cost sharing problem of [5]. To design competitive mechanisms for generalized auction problems that do not require the assumption of competition among the bidders in each market (e.g., allowing for singleton markets), global mechanisms must be employed. Even for simple special cases of the multicast tree, this problem is difficult. Special cases, e.g., segmented markets with no cost (for all $\mathbf{r}$, $c(\mathbf{r}) = 0$), that do permit natural global mechanisms are difficult to generalize.

Our formulation for generalized auctions captures a wide variety of allocation mechanisms; however, more general models can be considered. In particular, our model could be extended to allow cost functions that can take into account how many items are allocated to each market instead of just whether or not any goods were allocated to the market.

Acknowledgements

We gratefully acknowledge Gagan Aggarwal, Brian Babcock, Pedram Keyani, Doantam Phan and An Zhu for pointing out an error in an earlier version of this paper.

6. 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] A. Archer and E. Tardos. Frugal Path Mechanisms. In *To appear, SODA*, 2002.

[3] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proc. 13th Symp. on Discrete Alg.* ACM/SIAM, 2002.

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

[5] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. In *Proc. of 32nd Symposium Theory of Computing*, pages 218–226. ACM Press, New York, 2000.

[6] A. V. Goldberg, J. D. Hartline, 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, 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] J. Hershberger and S. Suri. Vickrey Pricing in network routing: Fast payment computation. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[10] K. Jain and V. Vazirani. Applications of approximation to cooperative games, 2001.

[11] P. Klemperer. Auction theory: A guide to the literature. In *Journal of Economic Surveys*, pages 227–286. 13(3), 1999.

[12] A. Mas-Collel, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.

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

[14] N. Nisan. Algorithms for Selfish Agents: Mechanism Design for Distributed Computation. In *16th Annual Symp. on Theoretical Aspects of Comp. Sci.*, 1999.

[15] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proc. of 31st Symposium on Theory of Computing*, pages 129–140. ACM Press, New York, 1999.

[16] N. Nisan and A. Ronen. Computationally feasible VCG Mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.

[17] C. Papadimitriou. Algorithms, Games and the Internet. In *Proc. of 33rd Symposium Theory of Computing*, pages 749–752. ACM Press, 2001.

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

[19] S. Shenker and V. Vazirani. In *Private Communication*, 2001.

[20] Combinatorial Auctions: A survey. S. DeVries and R. Vohra. In *Unpublished manuscript*, 2000.

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