Envy-Free Auctions for Digital Goods

Andrew V. Goldberg
Microsoft Research, SVC/5
1065 La Avenida,
Mountain View, CA 94043.
goldberg@microsoft.com

Jason D. Hartline
Computer Science Department
University of Washington
Seattle, WA 98195.
hartline@cs.washington.edu

ABSTRACT

We study auctions for a commodity in unlimited supply, e.g., a digital good. In particular we consider three desirable properties for auctions:

- *Competitive*: the auction achieves a constant fraction of the optimal revenue even on worst case inputs.
- *Truthful*: any bidder's best strategy is to bid the maximum value they are willing to pay.
- *Envy-free:* after the auction is run, no bidder would be happier with the outcome of another bidder (for unlimited supply auctions, this means that there is a single sale price and goods are allocated to all bidders willing to pay this price).

Our main result is to show that no constant-competitive truthful auction is envy-free. We consider two relaxations of these requirements, allowing the auction to be untruthful with vanishingly small probability, and allowing the auction to give non-envy-free outcomes with vanishingly small probability. Under both of these relaxations we get competitive auctions.

1. INTRODUCTION

Consider an auction for multiple identical items where each consumer desires exactly one item. A natural outcome of such an auction is the one defined by a sale price such that all consumers with bids above this price win and pay the price and all consumers with bids below the price lose. We call such outcomes *envy-free* because, for bids equal to consumer utilities, no consumer would prefer another consumer's outcome. *Envy-free auctions*,[1] auctions that always produce envy-free outcomes, are natural and desirable for consumer acceptance. While some auctions, such as the classical k-item Vickrey auction [12] and some optimal Bayesian auctions [3, 9] are envy-free, other auctions, in particular profit-maximizing competitive auctions [7], are not. In this paper we study compatibility of the envy-free property with the desirable auction properties of truthfulness and worst case profit maximization.

We consider the design of profit-maximizing auctions for digital goods such as downloadable audio files and pay-per-view television.[2] Since the marginal cost of duplicating a digital good is negligible and digital goods are freely disposable, we can assume that the auctioneer has an *unlimited supply* of items for sale. In this paper we are mostly interested in the good being sold is a *mass-market* commodity and that it is desirable to sell many copies of it. Although some of our results apply in a more general setting, thinking about the mass-market case gives the right context for our results. We assume that the bidders in the auction each have a private *utility value*, the maximum value they are willing to pay for the good. The auction takes as input bids from each of the consumers and determines which bidders receive a copy of the good and at what price. We assume that each consumer bids so as to maximize their own personal welfare, i.e., the difference between their utility value and the price they must pay for the good. We follow a standard approach from the field of mechanism design and study single round, sealed bid auctions that are *truthful*,[3] i.e., each bidder's personal welfare is maximized when they bid their true utility value.

When truthful mechanisms for a problem do not exist, it is natural to look at relaxations of the requirement that the mechanism be completely truthful. We adopt the notion from [1] that an auction is *truthful with probability* $1-\epsilon$ if the probability that any bidder can benefit from a untruthful bid is bounded by ϵ. An auction is *truthful with high probability* if ϵ tends to zero with some parameter in the input, e.g., the number of winners in an optimal auction. For other solution concepts related to approximate or probabilistic truthfulness, see for example [10, 11].

The previous work of [5, 6, 7] considered the design of truthful auction mechanisms for maximizing the profit of the seller under unknown market conditions. In contrast to the traditional approach from economics to profit maximization, which is to give an average case (Bayesian) analysis of an auction that is endowed with prior knowledge of the distribution from which the bidders utility values are drawn, these works analyze auctions in which the auctioneer has no

[1] We note that for the special case of auctions for a single commodity, the envy-free property is also referred to as *uniform-price*, *single-priced*, and *non-discriminatory*.

[2] In fact, these results apply to nonrival, or zero marginal cost, goods.

[3] Also known as *strategy-proof* or *incentive-compatible*

EC'03, June 9–12, 2003, San Diego, California, USA.

a priori knowledge of bidder utility values.

To analyze uninformed auctions, Goldberg et al. [7] proposed the use of worst case competitive analysis[4] wherein the performance of an uninformed auction is measured by comparing it to the performance of an optimal auction that is completely informed. In this setting they gave truthful auctions that are *competitive*, i.e., that always obtain a constant factor of the profit of an optimal informed auction. This work was extended in [5, 6] to improve the constant factor using new auction design and analysis techniques. Our analysis differs from that of Goldberg et al. in that we assume that we are in the mass-market case and that it is always desirable to sell many items.

Unfortunately, all known competitive auctions have the property that a bidder in the auction may get rejected while another bidder wins and pays a lower price than the first bidder's bid[5]. In this case, the first bidder would be envious of the second bidder's outcome. Thus, these auctions are not envy-free.

There is a good reason why none of the constant-competitive auctions considered to date have been envy-free. Our main result, presented in Section 3, is that it is not possible to design an auction that is constant competitive, truthful, and always has an envy-free outcome. We show that any auction that is always envy-free and truthful has a competitive ratio of $\Omega(\log n/\log\log n)$, where n is the number of bids. This bound is fairly tight: In Section 4 we give a truthful envy-free auction that is $O(\log n)$-competitive.

In light of this fundamental limitation, in Section 5 we consider two approaches to the relaxation of our goals. The first approach is to relax the envy-free requirement, requiring envy-free outcomes with high probability but not always. Another approach, recently employed by Archer et al. on a different auction mechanism design problem [1], is to relax the truthfulness requirement by requiring only that the auction be truthful with high probability. In both cases, the probability is over random coin tosses made by the randomized auction mechanism and not input. We present two auctions that are constant competitive, one always always truthful but only probabilisticly envy-free, another always envy-free but only probabilisticly truthful. Both of these auctions are based on the *Consensus Revenue Estimate* (CORE) technique of [6].

Auctions that do not have the envy-free property are impractical selling mechanisms for markets where consumers object to differential pricing [2]. For these applications, auctions that are truthful and almost envy-free or envy-free and almost truthful may be acceptable.

2. DEFINITIONS

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

Definition 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., [8]).*
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 a deterministic 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.

Finally, we formally define the notion of truthfulness.

Definition 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* $b_i = u_i$*. A randomized auction is truthful if it is a probability distribution over truthful deterministic auctions.*

We will be considering truthful auctions only until Section 5.2, where we define and discuss auctions that are truthful with high probability. As bidding u_i is a dominant strategy for bidder i in a truthful auction, for the remainder of this paper, we assume that $b_i = u_i$ (even for auctions that are only truthful with high probability).

The notion of a *bid-independent auction* formalizes the observation that in a truthful auction the price a bidder pays if they win should be independent of their bid value.

Definition 3. *[7] Let f be a function from bid vectors to prices (non-negative real numbers). The* deterministic bid-independent auction defined by f *is* $\mathcal{A}_f$*. For each bidder i:*

1. $t_i \leftarrow f(\mathbf{b}_{-i})$.
(where $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n)$*)*

[4] A technique borrowed from the analysis of on-line algorithms where the performance of an on-line algorithm is measured in comparison to that of an optimal off-line algorithm.

[5] In addition some of these auctions sell to different bidders at different prices; though, most of them can be adapted to have a single sale price for winners.

2. *If $t_i < b_i$, set $x_i \leftarrow 1$ and $p_i \leftarrow t_i$. (Bidder i wins.)*

3. *If $t_i > b_i$ set $x_i = p_i = 0$. (In this case, we say that bidder i is rejected.)*

4. *Otherwise, if $t_i = b_i$ the auction can either accept the bid at price t_i or reject it.*

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

For example, by setting $f = \max$ for all i and breaking ties arbitrarily, we obtain the 1-item Vickrey auction, i.e., 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 [12].

For auctions as defined above, we have the following result:

THEOREM 1. [5] *An auction is truthful if and only if it is equivalent to a bid-independent auction.*

We now formalize our notion of an envy-free auction.

DEFINITION 4. *The outcome of an auction is* envy-free *if there is a price, t, such that every winning bidder pays t, all bidders with bid values greater than t win, and all bidders with values lower than t lose. Bidders with bid values equal to t may either win or lose. Call such an outcome a* t-envy-free outcome.

From this definition it is clear that envy-free outcomes are completely specified by a number t and a description of whether tying bids, bids with $b_i = t$, win or lose. We note that our lower bounds will assume the more general case that tying bids can be treated differently. This is consistent with envy-free as such bidders have zero profit if they win and zero profit if they lose and are therefore assumed to have no preference over the two possible outcomes. For our upper bounds we will give auctions that satisfy the stronger condition that tying bidders are treated the same way, either they all win or they all lose. In all of our analyses it will be largely irrelevant how tying bids are treated; therefore, the relevant characteristic of the outcome is the value t.

We will be gauging the performance of a envy-free auction using competitive analysis. Following [7] we compare our auction profit to the profit of an optimal omniscient auction, one that has perfect knowledge of the bidder values in advance. As we are interested in the *mass-market* case, i.e., when the number of winners, m, in the optimal omniscient auction is large, we will be looking for a competitive ratio $\beta(m)$ to be a non-increasing function of m.

DEFINITION 5. *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 the vector $\mathbf{b}$. 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*

$$\mathcal{F}(\mathbf{b}) = \max_k kv_k.$$

In this paper we will be using the following notion of competitiveness.

DEFINITION 6. *We say that auction, $\mathcal{A}$, is $\beta(m)$-competitive for mass-markets if, for all bid vectors $\mathbf{b}$ such that $\mathcal{F}$ sells at least m items, we have*

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

As always in this paper, the expectation is over the randomized choices of the auction.

As shown in [7], if the optimal solution of $\mathcal{F}$ is to sell to only one bidder, then it may be impossible for an auction to be $\beta(1)$ competitive with $\beta(1)$ bounded by a constant. Thus, our goal will be to design an auction with low constant $\beta(2)$ and $\lim_{m\to\infty} \beta(m) = 1$ for $m > 2$. This framework is not as strong as the worst case framework of [5], which allows measurement of auction performance even on $\mathbf{b}$ where $\mathcal{F}$ sells only one item. While it is easy to modify our definition to generalize that of [5]; we do not do so here because the more general definition is not necessary for our analysis of the mass-market case.

However, the weaker framework allows a more precise characterization of auction performance in the more common case.

3. LOWER BOUND

In this section, we show that no envy-free auction can be $o(\log n/\log\log n)$ competitive.

For a particular input $\mathbf{b}$, an auction that always yields an envy-free outcome induces a probability distribution on values t such that the outcome is t-envy-free. Let $t(\mathbf{b})$ be the random variable for the price used by a auction with envy-free outcomes and define $p_{\mathbf{b}}(x)$ as

$$p_{\mathbf{b}}(x) = \mathbf{Pr}[t(\mathbf{b}) \leq x].$$

We begin by showing that any auction with envy-free outcomes has the property that the distribution of $t(\mathbf{b})$ on values less than x is independent of all bids with values above x.

LEMMA 2. *Let $\mathcal{A}$ be any truthful envy-free auction $\mathcal{A}$ with t and $p_{\mathbf{b}}$ defined as above. For all bids $\mathbf{b}$, $\mathbf{b}'$, values x, and subsets of bidders S such that all $i \in S$ have $b_i > x$ and $b'_i > x$, and $i \notin S$ have $b_i = b'_i$, $t(\mathbf{b})$ and $t(\mathbf{b}')$ have the same distribution on values at least x. That is,*

$$\forall y \leq x,\ p_{\mathbf{b}}(y) = p_{\mathbf{b}'}(y).$$

PROOF. Assume for a contradiction that there exists $\mathbf{b}$ and $\mathbf{b}'$ differing on some subset S (as defined above) such that $t(\mathbf{b})$ and $t(\mathbf{b}')$ do not have the same distribution on values at most x. Let $S = \{i_1, \ldots, i_k\}$ and $S_j = \{i_\ell \ :\ \ell \leq j\}$. Let $\mathbf{b}^{(j)}$ be equal to $\mathbf{b}$ except for $\mathbf{b}_{S_j} = \mathbf{b}'_{S_j}$. Note that $\mathbf{b}^{(0)} = \mathbf{b}$ and $\mathbf{b}^{(k)} = \mathbf{b}'$. It must be that for some j^* there exists $y \leq x$ such that $p_{\mathbf{b}^{(j^*-1)}}(y) \neq p_{\mathbf{b}^{(j^*)}}(y)$. However this violates truthfulness as the distribution of prices for bidder i_{j^*} when the bids are $\mathbf{b}^{(j^*)}$ is not bid-independent: if bidder i_{j^*} changes their bid to $b'_{i_{j^*}}$ they have a different distribution of prices. An auction that is not bid-independent is not truthful (Theorem 1). □

THEOREM 3. *No truthful envy-free auction is $\beta(m)$-competitive with $\beta(m) \in O(\log \frac{n}{m}/\log\log \frac{n}{m})$. In particular, if m is a constant, e.g., $m = 2$, we have $\beta(m) \in O(\log n/\log\log n)$.*

PROOF. For any integer $d \geq 2$, we show that no envy-free auction, $\mathcal{A}$, is $d/2$-competitive on an input $\mathbf{b}$ of size $n = md^d$ that has at least m winners. Since $d > \log(n/m)/\log\log(n/m)$, this gives the theorem. Assume for a contradiction that some auction, $\mathcal{A}$, is $d/2$-competitive.

Let $\mathbf{b}^*$ be the bid vector with $n/d^k - n/d^{k+1}$ bids at value n^k for $0 \leq k < d$ and m bids at n^d. Let $\mathbf{b}^{(k)}$ be identical to $\mathbf{b}^*$ except for the largest n/d^{k+1} bids, which are decreased to $n^k + \epsilon$:

$$b_i^{(k)} = \begin{cases} n^k + \epsilon & \text{if } i \leq n/d^{k+1} \\ b_i^* & \text{otherwise.} \end{cases}$$

There are two observations about these bid vectors that will be useful.

- $\mathcal{F}(\mathbf{b}^{(k)}) = n^{k+1}/d^k$ as the highest n/d^k bidders win at price n^k.
- $\mathcal{A}$ on $\mathbf{b}^{(k+1)}$ may place probability mass on prices at n^k or below. An upper bound on the contribution of this mass to the expected revenue of $\mathcal{A}$ is $\mathcal{F}(\mathbf{b}^{(k)}) = n^{k+1}/d^k$.

Define $t(\mathbf{b})$ and $p_{\mathbf{b}}(x)$ for $\mathcal{A}$ as above. We will show by induction that $\mathcal{A}$ has

$$p_{\mathbf{b}^{(k)}}(n^k) \geq k + 1/d.$$

This implies that for $\mathbf{b}^*$ we have

$$p_{\mathbf{b}^*}(n^{d-1}) = 1.$$

This is a contradiction, because if this were the case, $\mathcal{A}$ would have expected revenue at most n^d/d^{d-1} and could not be $2d$-competitive with $\mathcal{F}(\mathbf{b}^*) = mn^d$.

We will be using the fact that $\mathcal{A}$ is envy-free by invoking Lemma 2 in the inductive step to guarantee that $p_{\mathbf{b}^{(k+1)}}(n^k) = p_{\mathbf{b}^{(k)}}(n^k) \leq k + 1/d$.

For the base case, we show that $p_{\mathbf{b}^{(0)}}(1) \geq 1/d$. Note that $\mathbf{b}^{(0)}$ is defined as:

$$b_i^{(0)} = \begin{cases} 1 + \epsilon & \text{if } i \leq n/d \\ 1 & \text{otherwise.} \end{cases}$$

On this input $\mathcal{F}$ is n and as $d/2$-competitive auction, $\mathcal{A}$ must have expected revenue of at least $2n/d$. Suppose that $\mathcal{A}$ has $p_{\mathbf{b}^{(0)}}(1) = p$. Given this constraint, the best revenue is achieved by putting all probability mass p on 1 and the remaining probability mass $1 - p$ on $1 + \epsilon$. This gives the following bound on the expected revenue of $\mathcal{A}$:

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(0)})\Big] \leq np + (1-p)(1+\epsilon)(n/d).$$

This is an increasing function of $p \in [0,1]$. For $p = 1/d$ we have the following:

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(0)})\Big] \leq n/d + (1 - 1/d)(1+\epsilon)(n/d).$$

For ϵ chosen such that $(1+\epsilon)(1-1/d) = 1$ we have

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(0)})\Big] \leq 2n/d.$$

Therefore p must be at least $1/d$ for $\mathcal{A}$ to be $d/2$-competitive.

As the inductive step, assume that $p_{\mathbf{b}^{(k-1)}}(n^{k-1}) \geq k/d$ and consider running $\mathcal{A}$ on $\mathbf{b}^{(k)}$. By Lemma 2 and the definition of $\mathbf{b}^{(k)}$, $p_{\mathbf{b}^{(k-1)}}(n^{k-1}) = p_{\mathbf{b}^{(k)}}(n^k)$. Therefore, the probability mass remaining to be placed on values strictly larger than n^{k-1} is $P = 1 - p_{\mathbf{b}^{(k)}}(n^{k-1}) < 1 - k/d$. $\mathcal{A}$ must place p of this remaining mass on values at most n^k and the rest on higher values. Note that by definition, $\mathbf{b}^{(k)}$ has n/d^k bids at value strictly greater than n^{k-1}: it has n/d^{k+1} of them at $n^k + \epsilon$ and the remaining at n^k.

Thus, the most revenue can be obtained by placing all of p on n^k and all of the $P - p$ remaining mass on $n^k + \epsilon$. As discussed above, the expected revenue due to probability mass on values at most n^{k-1} is at most n^k/d^{k-1}. Thus, $\mathcal{A}$'s expected revenue is bounded as follows:

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(k)})\Big] \leq pn^k n/d^k + (P-p)(n^k+\epsilon)n/d^{k+1} + n^k/d^{k-1}$$

This is an increasing function of p. For $p = 1/d$ we have

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(k)})\Big] \leq n^k n/d^{k+1} + (P-p)(n^k+\epsilon)n/d^{k+1} + n^k/d^{k-1}.$$

Routine manipulation using the fact $d \geq 2$ shows that

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(k)})\Big] \leq 2n^{k+1}/d^{k+1}.$$

$\mathcal{F}$ of $\mathbf{b}^{(k)}$ is n^{k+1}/d^k, thus p is at least $1/d$. This gives the inductive claim that $p_{\mathbf{b}^{(k)}}(n^k) \geq k + 1/d$. $\square$

4. A LOG-COMPETITIVE, TRUTHFUL, ENVY-FREE AUCTION

We now give an auction that is truthful, envy-free, and has $\beta(m) \in \Theta(\log n)$ for all $m \geq 2$.

THEOREM 4. *There exists a truthful auction that is* $\Theta(\log n)$*-competitive.*

PROOF. First consider the k-item Vickrey auction that sells to the k highest bidders at the $k+1$st highest bid value. This auction is envy-free all winners pay the $k+1$st price and all losers bid at most the $k+1$st price.[6]

Our $(2\log n)$-competitive auction picks i uniformly at random from $[0, \ldots, \lfloor \log n \rfloor]$ and runs the 2^i-item Vickrey auction. The worst-case value of $\beta(m)$ for this auction occurs on $\mathbf{b}$ with $n-m$ bids at value 0 and m bids at distinct values in $[v, v+\epsilon]$ for any positive v and ϵ. Note that the 2^i-Vickrey auction gets revenue approximately $2^i v$ if $2^i < m$ and zero otherwise. Since we choose each auction with probability $1/\log n$, our expected revenue on $\mathbf{b}$ (with m winners in $\mathcal{F}$) is at least:

$$\begin{aligned} \mathbf{E}[\mathcal{R}] &= \tfrac{v}{\log n} \sum_{i=0}^{\lceil \log m \rceil - 1} 2^i = \tfrac{v}{\log n}\left(2^{\lceil \log m \rceil} - 1\right) \\ &\geq \tfrac{v}{\log n}(m-1) \geq \tfrac{m-1}{m \log n}\mathcal{F}(\mathbf{b}) \geq \tfrac{\mathcal{F}(\mathbf{b})}{2\log n}. \end{aligned}$$

$\square$

[6] In place of k-Vickrey we can use a slight modification of k-Vickrey that does not arbitrarily break ties in the case that there are more than one bid value tied with the $k+1$st price. This version of Vickrey may sell more than k items. For example, consider the auction for $k = 2$ on $(4, 2, 2, 1)$. The the first three bidders win the auction and pay 2 each.

5. RELAXATIONS

In Section 3, we showed that no truthful, envy-free auction can be better than $\log n / \log\log n$-competitive. Nonetheless, envy-free outcomes are still a desirable goal. In this section, we present two auctions that come close to being competitive, truthful, and envy-free at the same time. The first auction is truthful, but is only envy-free with high probability; the second auction is envy-free, but only truthful with high probability.

5.1 CORE and envy-free outcomes

In this section, we relax the condition that the auction outcome is always envy-free and instead consider auctions that are truthful but only envy-free with high probability. One such auction is the Consensus Revenue Estimate (CORE) auction from [6]. We review the mass-market version of the CORE auction and discuss its properties.

The CORE auction combines two general ideas that have proven to be successful for designing competitive mechanisms [4, 5]. The first is that of a *profit extractor*. A profit extractor is a truthful parameterized mechanism that, given a target profit, extracts that profit from the bidders if it is possible. For basic auctions the *maximum extractable profit* is the same as the profit of the optimal omniscient auction $\mathcal{F}$. The profit extraction mechanism for basic auctions is the following special case of the Moulin-Shenker cost sharing mechanism [8]:

> CostShare_R: Given bids $\mathbf{b}$, find the largest k such that the highest k bidders can equally share the cost R. Charge each of these bidders R/k. If no k exists, no bidders win.

This mechanism has the following properties:

- It is truthful [8].
- It has profit R if $R \leq \mathcal{F}$ (otherwise there are no winners).
- It is envy-free.

To see the last point, note that by definition all winners pay the same price R/k and furthermore, all other bidders bid less than R/k (otherwise a larger set of bidders could share the cost evenly).

The second technique for the design of competitive mechanisms is that of using a bid-independent *consensus estimate*. Note that Theorem 1 and the truthfulness of CostShare_R imply that it is implemented by some bid-independent function, cs_R. That is, the bid-independent auction defined by $f(\mathbf{b}_{-i}) = \text{cs}_R(\mathbf{b}_{-i})$ is exactly CostShare_R. Consider the auction, $\mathcal{A}$, parameterized by function $r(\cdot)$ that is defined by bid-independent function $f(\mathbf{b}_{-i}) = \text{cs}_{r(\mathbf{b}_{-i})}(\mathbf{b}_{-i})$. Note that if r is *a consensus*, i.e., $r(\mathbf{b}_{-i}) = R$ for all i, then $\mathcal{A}$ is identically CostShare_R.

Recall that CostShare_R gives profit R in the case that the profit R is achievable, i.e., $R \leq \mathcal{F}$. The consensus estimate technique is to construct an $r(\cdot)$ that gives an estimate of $\mathcal{F}$ and is constant on $\mathbf{b}_{-i}$ for all i (i.e., it is a consensus). For the case that $\mathcal{F}(\mathbf{b}_{-i})$ is a constant fraction of $\mathcal{F}(\mathbf{b})$, i.e.,

$$\tfrac{1}{\rho}\mathcal{F}(\mathbf{b}) \leq \mathcal{F}(\mathbf{b}_{-i}) \leq \mathcal{F}(\mathbf{b}),$$

it is possible to pick an $r(\cdot)$ from a distribution of functions that are good estimates of $\mathcal{F}$ such that with high probability (in the choice of r), r is a consensus.

Parameterized by constant c with $c > \rho$, we choose $r(\cdot)$ as follows. For U be uniformly distributed on $[0, 1]$, define $r(\mathbf{b})$ is $\mathcal{F}(\mathbf{b})$ rounded down to nearest c^{i+u} for integer i.

LEMMA 5. *[6] For $\rho < c$ and $\mathbf{b}$ with $\mathcal{F}(\mathbf{b}_{-i}) \in [\mathcal{F}(\mathbf{b})/\rho, \mathcal{F}(\mathbf{b})]$, the probability that r is a consensus is $1 - \log_c \rho$.*

Combining the *consensus estimate* solution with the *profit extractor*, CostShare_R, we get the following auction:

DEFINITION 7 (CORE_c). *For constant c,*

- *pick U uniformly at random from $[0, 1]$, and*
- *let function $r(\cdot)$ be $\mathcal{F}(\cdot)$ rounded down to nearest c^{i+U} for integer i.*

The CORE_c auction is defined by bid-independent function $f(\mathbf{b}_{-i}) = \text{cs}_{r(\mathbf{b}_{-i})}(\mathbf{b}_{-i})$.

We say that CORE_c achieves a consensus at value R if $r(\mathbf{b}_{-i}) = R$ for all i. In this case, by definition, CORE_c's outcome is identical to that of CostShare_R. Thus, this outcome of CORE_c is envy-free.

LEMMA 6. *If CORE_c achieves consensus then its outcome is envy-free.*

The auction CORE_c is constant-competitive:

LEMMA 7. *[6] For $\rho < c$ and $\mathbf{b}$ with $\mathcal{F}(\mathbf{b}_{-i}) \in [\mathcal{F}(\mathbf{b})/\rho, \mathcal{F}(\mathbf{b})]$, the expected revenue of CORE_c is*

$$\frac{\mathcal{F}(\mathbf{b})}{\ln c}\left(\frac{1}{\rho} - \frac{1}{c}\right).$$

For mass-market applications, removing one bid does not change the maximum extractable profit by much. If the optimal mechanism, $\mathcal{F}$, sells m items then removing any bid can change the optimal by at most a factor of $1/\rho = (m-1)/m$. Plugging this into Lemma 5 and using the fact that the Taylor expansion of $\log(1+1/x) = \Theta(1/x)$, we have:

COROLLARY 8. *On $\mathbf{b}$ with $\mathcal{F}$ selling $m \geq 2$ items, CORE_c is envy-free (i.e., achieves a consensus) with probability at least*

$$1 - \log_c \tfrac{m}{m-1} = 1 - \Theta(1/m).$$

COROLLARY 9. *When $\mathcal{F}$ sells $m \geq 2$ items, the CORE_c auction is $\beta(m)$-competitive with*

$$\beta(m) = \frac{\ln c}{1 - \frac{1}{c} - \frac{1}{m}}.$$

In the limit CORE_c is $\lim_{m\to\infty} \beta(m) = c \ln c/(c-1)$ competitive.

The choice of c in the CORE_c auction is crucial. In order for consensus to work on bids $\mathbf{b}$ on which $\mathcal{F}$ sells at least m items, we need $c > m/(m-1)$. Otherwise there exists a set of bids such that no consensus is achievable. If we set $c = 1 + 1/\sqrt{m}$, then $\beta(m) \to 1$ as $m \to \infty$.

As noted in [6], in the case that CORE does not achieve a consensus, there are two sale prices, and the auction is not envy-free. For mass markets, such as television broadcasts where the number of viewers is in hundreds of thousands, the probability of a non-envy-free outcome is very small.

5.2 Truthful with high probability

Another approach to deal with the non-existence of truthful competitive auctions that always have envy-free outcomes is to relax the requirement that the auction always be truthful. Next we define a *probabilistically truthful* mechanism. Let m be the number of winners in the optimal auction on a given set of bids. We will be looking for a mechanism with good probabilistically truthful properties in terms of m.

Definition 8. *[1] An auction is* truthful with probability $1-\epsilon$ *if the probability that any bidder can benefit from an untruthful bid is at most* ϵ. *An auction is* truthful with high probability *if* $\epsilon \to 0$ *as* $m \to \infty$*, where* m *is the number of winners in the optimal auction* $\mathcal{F}$.

Definition 9 (CORE$'_c$). *For constant* c *and input* $\mathbf{b}$*, CORE$'_c$ is:*

1. *Pick* U *uniformly at random from* $[0,1]$.
2. *Let function* $r(\cdot)$ *be* $\mathcal{F}(\cdot)$ *rounded down to nearest* c^{i+U} *for integer* i.
3. *Run* $\text{CostShare}_{r(\mathbf{b})}$ *on* $\mathbf{b}$.

Lemma 10. *For bids* $\mathbf{b}$ *and choice of* U *fixed identically for both CORE$_c$ and CORE$'_c$, if CORE$_c$ is a consensus then CORE$'_c$ is truthful.*

Proof. Let $R = r(\mathbf{b})$. To prove the lemma, we must show that for U such that CORE$_c$ is a consensus, no bidder can benefit from bidding any value other than their true utility value. First we note that CORE$'_c$ runs $\text{CostShare}_{r(\mathbf{b})}$ on $\mathbf{b}$. Consider the effect of bidder i changing their bid to b'_i resulting in bid vector $\mathbf{b}'$. Bidder i cannot benefit from bidding any b'_i such $r(\mathbf{b}') = r(\mathbf{b}) = R$ because CostShare_R is truthful and therefore bidder i's best strategy in mechanism CostShare_R is to bid b_i.

The fact that CORE$_c$ is a consensus for this value of U means that $r(\mathbf{b}_{-i}) = r(\mathbf{b}) = R$ for all i. Again consider $\mathbf{b}'$ identical to $\mathbf{b}$ except for bidder i bidding b'_i. For $b'_i \in [0, b_i]$, since $r(\cdot)$ is a monotonically increasing function of the bids, $r(\mathbf{b}_{-i}) = r(\mathbf{b}') = r(\mathbf{b}) = R$. Thus, no bidder can lower the the value of $r(\mathbf{b})$.

We now consider bidder i raising their bid enough to make $r(\mathbf{b}') > r(\mathbf{b})$. Although we do not go into the details here, it is not difficult to show that $\text{cs}_R(\mathbf{b}_{-i})$ (the bid-independent function implementing CostShare_R), for $\mathbf{b}_{-i}$ fixed, is a monotone increasing function of R. Thus, if bidder i raises their bid to raise R, the price offered them in the by $\text{cs}_{R'}$ is going to be higher than for cs_R and therefore, bidder i would be worse off. □

The following corollary follows from the fact above lemma and Corollary 8.

Corollary 11. *On* $\mathbf{b}$ *with* $\mathcal{F}$ *selling* $m \geq 2$ *items, CORE$'_c$ is truthful with probability*

$$1 - \log_c \tfrac{m}{m-1} = 1 - \Theta(1/m).$$

We now consider the performance of CORE$'_c$. We can view the use of $r(\cdot)$ in CORE$'_c$ as a consensus problem with $\rho = 1$ since the same value, $r(\mathbf{b})$, is used for all bidders. This allows us to make use of Lemma 7 directly to obtain the following corollary.

Corollary 12. *The expected revenue of CORE$'_c$ is*

$$\frac{\mathcal{F}(\mathbf{b})}{\ln c}\left(1 - \frac{1}{c}\right).$$

Conclusions and Extensions

The CORE technique of computing a consensus revenue estimate and using a profit extractor to obtain the estimated revenue has several notable properties. First note that for the mass-market sale of a good, digital or otherwise, we would expect that the result of the sale should not depend on the actions of any one bidder. In this respect, the outcome of the CORE auction fits with our expectations: with high probability the results of the CORE auction are not affected by any one bidder.

The second observation is that, with high probability, the CORE auction is equivalent to the profit extraction mechanism that it is based on, which is a special case of the Moulin-Shenker [8] cost sharing mechanism. We exploited the fact that the latter mechanism is envy-free. The mechanism is also *group strategy-proof*: no coalition of bidders can collude by bidding untruthfully so as some members of the coalition strictly benefit without causing other members of the coalition to be strictly worse off. As such, both CORE$_c$ and CORE$'_c$ have collusion resistant properties. We are currently formalizing the extent of these properties.

In this paper we consider truthfulness with high probability to get an auction that is both envy-free and competitive. In our analysis we make the assumption that bidders will still reveal their true utility values even though the mechanism is only truthful with high probability. In the case where each bidder is perfectly informed as to the bidding strategies and utility values of other bidders, any non-truthful auction would fail to obtain true bids, as each bidder could calculate their own optimum bid. However, in the presence of uncertainty about the strategies or utility values, this calculation is no longer straight-forward. A bidder in an auction that is truthful with high probability is faced with a choice. The bidder can bid truthfully, which with high probability is an optimal strategy. Alternatively, the bidder could try to gain by manipulating their bid on the basis of available information. If information is incomplete, the latter may be impossible.

More generally, we have shown that one can get stronger results by relaxing the notion of truthfulness. Relaxed truthfulness is an interesting direction for future research. For example, is it possible to get better competitive ratios under a reasonable relaxation? Which relaxations are reasonable?

Acknowledgments

We are grateful to Amos Fiat for suggesting that we consider envy-free outcomes and for his contributions to the initial discussions on the topic; to Joan Feigenbaum for pointing us to many helpful references for approximate notions of truthfulness; and to Anna Karlin for many helpful discussions. We also would like to than anonymous referees for their comments.

6. REFERENCES

[1] A. Archer, C. Papadimitriou, K. Tawar, and E. Tardos. An Approximake Truthful Mechanism for Combinatorial Auctions with Single Parameter

Agents. In *Proc. 14th Symp. on Discrete Alg.* ACM/SIAM, 2003.

[2] Associated Press. Amazon may spell end for 'dynamic' pricing. *USA Today*, 29 September 2000. Available at: `http://www.usatoday.com/life/cyber/tech/cti595.htm`.

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

[4] K. Deshmukh, A. Goldberg, J. Hartline, and A. Karlin. Truthful and Competitive Double Auctions. In *Proc. 10th European Symposium on Algorithms*. Springer–Verlag, 2002.

[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 and J. D. Hartline. Competitiveness via Concensus. In *Proc. 14th Symp. on Discrete Alg.* ACM/SIAM, 2003.

[7] 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.

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

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

[10] R. Radner. Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives. *J. Economic Theory*, 22:136–154, 1980.

[11] J. Schummer. Almost dominant strategy implementation. Technical report, MEDS, Kellogg Graduate School of Management, 2001. Available from `http://www.kellogg.nwu.edu/faculty/schummer/ftp/research/esp.ps`.

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