

# Truthful and Competitive Double Auctions

Kaustubh Deshmukh<sup>1</sup>, Andrew V. Goldberg<sup>2</sup>,  
Jason D. Hartline<sup>1</sup>, and Anna R. Karlin<sup>1</sup>

<sup>1</sup> Computer Science Department, University of Washington.

{kd,hartline,karlin}@cs.washington.edu

<sup>2</sup> Microsoft Research, 1065 La Avenida, SVC 5, Mountain View, CA 94043.

goldberg@microsoft.com

**Abstract** In this paper we consider the problem of designing a mechanism for double auctions where bidders each bid to buy or sell one unit of a single commodity. We assume that each bidder's utility value for the item is private to them and we focus on truthful mechanisms, ones where the bidders' optimal strategy is to bid their true utility. The profit of the auctioneer is the difference between the total payments from buyers and total to the sellers. We aim to maximize this profit. We extend the competitive analysis framework of basic auctions [9] and give an upper bound on the profit of any truthful double auction. We then reduce the competitive double auction problem to basic auctions by showing that any competitive basic auction can be converted into a competitive double auction with a competitive ratio of twice that of the basic auction. In addition, we show that better competitive ratios can be obtained by directly adapting basic auction techniques to the double auction problem. This result provides insight into the design of profit maximizing mechanisms in general.

## 1 Introduction

Dynamic pricing mechanisms, and specifically auctions with multiple buyers and sellers, are becoming increasingly popular in electronic commerce. We consider *double auctions* in which there is one commodity in the market with multiple buyers and sellers each submitting a single bid to either buy or sell one unit of the commodity (E.g. [7]). The numerous applications of double auctions in electronic commerce, including stock exchanges, business-to-business commerce, bandwidth allocation, etc. have led to a great deal of interest in fast and effective algorithms [22,24].

For double auctions, the auctioneer, acting as a broker, is faced with the task of matching up a subset of the buyers with a subset of the sellers of the same size. The auctioneer decides on a price to be paid to each seller and received from each buyer in exchange for the transfer of one item from each of the selected sellers to each of the selected buyers. The *profit of the auctioneer* is the difference between the prices paid by the buyers and the prices paid to the sellers.

We assume that each bidder has a private utility value for the item. For the buyers this utility value is the most that they are willing to pay for the item, and

and for the sellers it is the least they are willing to sell for. We focus on double auction mechanisms that are *truthful*: the best strategy of a selfish bidder that is attempting to maximize their own gain is to bid their true utility value.

The traditional economics approach to the study of profit maximizing auctions is to construct the optimal Bayesian auction given the prior distribution from which the bidders' utility values are drawn (e.g., [3,17]). In contrast, following [10,6,9], we attempt to design mechanisms that maximize profit under *any* market conditions. As in competitive analysis of online algorithms, we gauge a truthful double auction mechanism's performance on a particular bid set by comparing it against the profit that would be achieved by an "optimal" auction,  $\mathcal{M}_{\text{opt}}$ , on the same bidders.

If, for every bid set, a particular truthful double auction mechanism  $\mathcal{M}$  achieves a profit that is close to that of the optimal  $\mathcal{M}_{\text{opt}}$ , we say that the auction mechanism  $\mathcal{M}$  is *competitive* against  $\mathcal{M}_{\text{opt}}$ , or simply competitive. For example, we might be interested in constructing double auctions that are competitive with the optimal single-price omniscient mechanism,  $\mathcal{F}$ . This is the mechanism which, based on perfect knowledge of buyer and seller utilities, selects a single price  $b_{\text{opt}}$  for the buyers and a single price  $s_{\text{opt}}$  for the sellers. It then finds the largest  $k$  such that the highest  $k$  buyers each bid at least  $b_{\text{opt}}$  and the lowest  $k$  sellers each bid at most  $s_{\text{opt}}$ . It then matches these buyers and sellers up, paying all the sellers  $s_{\text{opt}}$  and charging each of the buyers  $b_{\text{opt}}$ . The profit of the auctioneer is thus  $k(b_{\text{opt}} - s_{\text{opt}})$ .

## 1.1 Results

This paper makes the following contributions:

- We extend the framework for competitive analysis of basic auctions to double auctions. This framework is motivated by a number of results bounding truthful mechanism. In particular, we show that no monotone<sup>3</sup> double auction (even a *multi-priced* mechanism) can achieve a higher profit than twice the optimal *single-priced* mechanism  $\mathcal{F}$  discussed above.
- We present a reduction from double auctions to basic auctions by showing how to construct a competitive double auction from any competitive basic auction while losing only a factor of two in competitive ratio.
- We show how the basic auction from [8] can be adapted to the double auction problem yielding better competitive ratio than one get by applying the aforementioned reduction. We also discuss the possibility of making similar adaptations for other profit maximizing mechanism design problems.

## 1.2 Related Work

We study profit maximizing single round double auctions when the utility value of each bidder is private and must be truthfully elicited. When the utilities are

---

<sup>3</sup> See Section 2.2 for the definition of monotone.

public values this problem becomes trivial. The following variants of the problem have been previously studied.

When the goal is not to maximize profit of the auctioneer, but to find an outcome that maximizes the *common welfare*, i.e., the sum of the profits of each of the bidders, subject to the constraint that the auctioneer's profit is non-negative, McAfee [15] gives a truthful mechanism that approaches optimal as the number of sold items in the optimal solution grows. Note that the Vickrey-Clarke-Groves [4,11,23] mechanism, the only mechanism that always gives the outcome that maximizes the common welfare, always gives a non-positive profit to the auctioneer (assuming *voluntary participation*<sup>4</sup>).

Our results are closely related to the *basic auctions* for a single item available in unlimited supply, e.g., for digital goods [10,9]. As such, the approach we take in this paper closely parallels that in [9]. Furthermore, as we explain later, the basic auction is a special case of the double auction where all sellers have utility zero.

An "online" version of the double auction, where bids arrive and expire at different times, was considered by Blum, Sandholm, and Zinkevich [2] (also known as a *continuous double auction* [24]). Their mechanism must make decisions without knowing what bids will arrive in the future. They consider the goals of optimizing the profit of the auctioneer and of maximizing the number of items sold. Their solution assumes that bidders are compelled to bid their true utility value despite the fact that the algorithms they develop are not truthful, i.e., the utility values of the bidders are public. An interesting open question left by our work is the problem of a profit maximizing online double auction for the private value model. For private values, an online variant of the basic auction problem was considered by Bar-Yossef et al. [1] in a competitive framework for profit maximization.

More generally, there has been a great deal of recent work at the intersection of game theory, economic theory and theoretical computer science [18,21]. 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., [14], chapter 23). One of the most important positive results in this field is aforementioned family of Vickrey-Clarke-Groves mechanisms.

Recent work in computer science pertaining to these fields has focused largely on merging the considerations of incentives (e.g. truthfulness) with considerations of computational complexity [12,18,19,20]. One of the first examples of such work is the paper of Nisan and Ronen [19] where the mechanism design framework is applied to some standard optimization problems in computer science, such as shortest paths and scheduling on unrelated machines.

Of course, auctions, be they traditional or combinatorial, have received a great deal of attention (see e.g., the surveys [5,13]).

---

<sup>4</sup> Defined in Section 2.

## 2 Preliminaries

We consider single-round, sealed-bid double auction mechanisms in which a set of bidders each wants to either buy or sell one out of a set of identical items. Bidders submit one sealed bid each, and publicly declare themselves to be either *buyers* or *sellers*. For buyers, the bid represents the maximum amount they are willing to pay for an item, whereas for sellers, the bid represents the minimum amount they are willing to sell the item for.

We denote by  $\mathbf{b}$  the vector of bid values associated with buyers and by  $\mathbf{s}$  the vector of all bid values associated with sellers. The  $i$ th component of  $\mathbf{b}$  (resp.  $\mathbf{s}$ ) is  $b_i$  (resp.  $s_i$ ), the bid value submitted by the  $i$ th buyer (resp. seller). We assume that the number of buyers is equal to the number of sellers, and we use  $n$  to denote this number.

Throughout the paper we will be using the notation  $b_{(i)}$  to represent the  $i$ th largest buyer bid and  $s_{(i)}$  for the  $i$ th smallest seller bid.

**Definition 1.** A single-round sealed-bid double-auction mechanism is one where:

- Given the two bid vectors  $\mathbf{b} = (b_1, \dots, b_n)$  and  $\mathbf{s} = (s_1, \dots, s_n)$ , the mechanism computes a pair of allocation vectors,  $\mathbf{x}$  and  $\mathbf{y} \in \{0, 1\}^n$ , and payment vectors  $\mathbf{p}$  and  $\mathbf{q} \in \mathbb{R}^n$ , subject to the constraints that:
  - The number of winning buyers is equal to the number of winning sellers, i.e.,  $\sum_i x_i = \sum_i y_i$ .<sup>5</sup>
  - $0 \leq p_i \leq b_i$  (resp.  $s_i \leq q_i$ ) for all winning buyers (resp. sellers) and that  $p_i = 0$  (resp.  $q_i = 0$ ) for all losing buyers (resp. sellers). These are the standard assumptions of no positive transfers and voluntary participation. See, e.g., [16].
- If  $x_i = 1$  buyer  $i$  wins (i.e., receives the item) and pays price  $p_i$ , otherwise we say that buyer  $i$  loses. If  $y_i = 1$  seller  $i$  wins (i.e., sells the item) and gets paid  $q_i$ , otherwise we say that seller  $i$  loses.
- The profit of the mechanism is  $\mathcal{M}(\mathbf{b}, \mathbf{s}) = \sum_i p_i - \sum_i q_i$ .

Note that the *basic auction* problem of [10] can be viewed as a special case of the double auction problem with all sell bids equal to zero.

We say the *mechanism is randomized* if the procedure used to compute the allocations and prices is randomized. Otherwise, the mechanism is *deterministic*. Note that if the mechanism is randomized, the profit of the mechanism, the output prices, and the allocation are random variables.

We use the following private value model for the bidders:

- Each bidder has a private utility value for the item. We denote the utility value for buyer  $i$  by  $u_i$  and the utility value for seller  $i$  by  $v_i$ .

---

<sup>5</sup> We assume that the auctioneer neither has any items for sale nor is willing to purchase any. For this reason, we can also assume that the number of buyer bids equals the number of seller bids. If there are any extra buyers or sellers, the auctioneer can earn the same amount of profit by ignoring the extra low bidding buyers or high bidding sellers.

- Each bidder bids so as to maximize their *profit*: For buyers (resp. sellers) this means they bid  $b_i$  (resp.  $s_i$ ) to maximize profit given by  $u_i x_i - p_i$  (resp.  $q_i - v_i y_i$ ).
- Bidders bid with full knowledge of the auctioneer’s strategy. However, the bidding occurs in advance (i.e., without knowledge) of any coin tossing done by a randomized auctions.
- Bidders do not collude.

Finally, we formally define the notion of truthfulness.

**Definition 2.** *We say that a deterministic double 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  $b_i = u_i$  for buyers and by setting  $s_i = v_i$  for sellers.*

**Definition 3.** *We say that a randomized auction is truthful if it can be described as a probability distribution over deterministic truthful auctions.*

As bidding  $u_i$  (resp.  $v_i$ ) is a dominant strategy for buyer  $i$  (resp. seller  $i$ ) in a truthful auction, in the remainder of this paper, we assume that  $b_i = u_i$  and  $s_i = v_i$  unless mentioned otherwise.

As an example of a truthful double auction, consider the natural generalization of the  $k$ -item Vickrey auction [23] to the double auction problem. This mechanism simply buys from the lowest  $k$  sellers at the  $(k + 1)$ -st lowest sell bid, and sells these items to the highest  $k$  buyers at the  $(k + 1)$ -st highest buy bid. It is easy to see that for any fixed  $k$ , this mechanism is truthful. However, it does not in general maximize profit.

## 2.1 Bid Independence

We describe a useful characterization of truthful mechanisms using the notion of *bid independence*. Let  $\mathbf{b}_{-i}$  denote the vector of bids  $\mathbf{b}$  with  $b_i$  removed, i.e.,  $\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, ?, b_{i+1}, \dots, b_n)$ . We call such a vector *masked*. Similarly, let  $\mathbf{s}_{-i}$  denote the masked vector of bids  $\mathbf{s}$  with  $s_i$  removed, i.e.,  $\mathbf{s}_{-i} = (s_1, \dots, s_{i-1}, ?, s_{i+1}, \dots, s_n)$ . Given bid vectors  $\mathbf{b}$  and  $\mathbf{s}$ , the bid-independent mechanism defined by the randomized functions  $f$  and  $g$  is defined as follows:

For each buyer  $i$ , if  $f(\mathbf{b}_{-i}, \mathbf{s}) \leq b_i$ , buyer  $i$  wins at the price  $p_i = f(\mathbf{b}_{-i}, \mathbf{s})$ . Otherwise, buyer  $i$  loses and makes no payment. Similarly for each seller  $i$ , if  $s_i \leq g(\mathbf{b}, \mathbf{s}_{-i})$ , seller  $i$  wins and receives a payment of  $q_i = g(\mathbf{b}, \mathbf{s}_{-i})$ . Otherwise, seller  $i$  loses and receives no compensation.<sup>6</sup>

The following theorem, which is a straightforward generalization of the equivalent result for basic auctions in [9], relates bid independence to truthfulness.

**Theorem 1.** *A double auction is truthful if and only if it is bid-independent.*

<sup>6</sup> In fact, bid-independence allows the inequalities,  $f(\mathbf{b}_{-i}, \mathbf{s}) \leq b_i$  and  $s_i \leq g(\mathbf{b}, \mathbf{s}_{-i})$  to be either strict or non-strict at the discretion of the functions  $f$  and  $g$ .

## 2.2 Monotonicity

We define the notion of monotone double auctions to characterize “reasonable” truthful mechanisms. Using standard terminology, we say that random variable  $X$  *dominates* random variable  $Y$  if for all  $x$

$$\Pr[X \geq x] \geq \Pr[Y \geq x].$$

**Definition 4.** *A double auction mechanism is monotone if it is defined by a pair of bid-independent functions  $f$  and  $g$  (each taking as input a buy vector and a sell vector, where one of the two vectors is masked) such that for any buy and sell vectors  $\mathbf{b}$  and  $\mathbf{s}$ , we have:*

- *For any pair of buyers  $i$  and  $j$  such that  $b_i \leq b_j$  the random variable  $f(\mathbf{b}_{-i}, \mathbf{s})$  dominates the random variable  $f(\mathbf{b}_{-j}, \mathbf{s})$ .*
- *For any pair of sellers  $i$  and  $j$  such that  $s_i \leq s_j$  the random variable  $g(\mathbf{b}, \mathbf{s}_{-i})$  dominates the random variable  $g(\mathbf{b}, \mathbf{s}_{-j})$ .*

To get a feel for this definition, observe that when  $b_i \leq b_j$  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 smaller bid  $b_i$  is visible in  $\mathbf{b}_{-j}$  whereas the larger bid  $b_j$  is visible in  $\mathbf{b}_{-i}$ . Intuitively, monotonicity means that if buyer bids are increased while keeping the seller bid vector constant, then the threshold prices output by the bid-independent function  $f$  increase. Similarly for the sellers.

## 2.3 Single Price Omniscient Mechanism

A key question is how to evaluate the performance of mechanisms with respect to the goal of profit maximization. A natural objective would be to design a mechanism that, on all inputs, achieves a profit close to that of the *optimal multiple-price omniscient mechanism*  $\mathcal{T}$  that requests each buyer to pay their utility value and pays to each seller their utility value (as long as positive profit can be made). Note that the profit of this mechanism  $\mathcal{T}(\mathbf{b}, \mathbf{s}) = \sum_i b_i - \sum_i s_i$ , if all buyers bid above all sellers.

**Definition 5.** *The optimal single price omniscient mechanism,  $\mathcal{F}$ , is the mechanism that uses the optimal single buy price and single sell price. It achieves the optimal single price profit of*

$$\mathcal{F}(\mathbf{b}, \mathbf{s}) = \max_i (b_{(i)} - s_{(i)}).$$

The facts that no reasonable truthful mechanism can achieve profit above  $2\mathcal{F}(\mathbf{b}, \mathbf{s})$  (as we will prove shortly) and  $\mathcal{F}(\mathbf{b}, \mathbf{s})$  can be  $\Theta(\mathcal{T}(\mathbf{b}, \mathbf{s})/\log n)$  (see [9]) imply that we cannot hope to be competitive with the optimal multi-price mechanism  $\mathcal{T}$ . This motivates using  $\mathcal{F}$  as a performance metric. Unfortunately, it is impossible to be competitive with  $\mathcal{F}$ . This is shown in [9] for basic auctions, a special case of the double auction, when the  $\mathcal{F}$  sells to only the highest bidder. Hence, we compare the performance of truthful mechanisms with the profit of the optimal single price omniscient mechanism that transfers at least two items from sellers to buyers.

**Definition 6.** The optimal fixed price mechanism that transfers at least two items,  $\mathcal{F}^{(2)}$ , has profit

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) = \max_{i \geq 2} i(b_{(i)} - s_{(i)}).$$

## 2.4 Competitive Mechanisms

We now formalize the notion of a competitive mechanism:

**Definition 7.** We say that a truthful double auction  $\mathcal{M}$  is  $\beta$ -competitive against  $\mathcal{F}^{(2)}$  if, for all bid vectors  $\mathbf{b}$  and  $\mathbf{s}$  the expected profit of  $\mathcal{M}$  satisfies

$$\mathbf{E}[\mathcal{M}(\mathbf{b}, \mathbf{s})] \geq \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})/\beta.$$

We say that  $\mathcal{M}$  is competitive if  $\mathcal{M}$  is  $\beta$ -competitive for some constant  $\beta$ .

## 3 Upper Bound on the Profit of Truthful Mechanisms

In this section, we show that the profit for all monotone double auction mechanisms is bounded by  $2\mathcal{F}(\mathbf{b}, \mathbf{s})$ . Goldberg et al. showed that for basic auctions this result holds without the factor of two:

**Theorem 2.** For input bids  $\mathbf{b}$ , no truthful monotone basic auction has expected profit more than  $\mathcal{F}(\mathbf{b})$  [9].

We conjecture that this bound holds for double auctions as well, though what we prove below is a factor of two worse.

**Lemma 1.** For any value  $v$  and buy and sell bids  $\mathbf{b}$  and  $\mathbf{s}$ , define  $\mathbf{b}'$  and  $\mathbf{s}'$  as  $b'_i = b_i - v$  and  $s'_i = v - s_i$  for  $1 \leq i \leq n$ . Then for any monotone double auction,  $\mathcal{M}$ :

$$\mathbf{E}[\mathcal{M}(\mathbf{b}, \mathbf{s})] \leq \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}').$$

*Proof.* Let  $\mathbf{x}$ ,  $\mathbf{y}$ ,  $\mathbf{p}$ , and  $\mathbf{q}$  be the outcome and prices when  $\mathcal{M}$  is run on  $\mathbf{b}$  and  $\mathbf{s}$ . Let  $X = \{i : x_i = 1\}$  and  $Y = \{i : y_i = 1\}$ . Note  $|X| = |Y|$ . Thus,

$$\begin{aligned} \mathcal{M}(\mathbf{b}, \mathbf{s}) &= \sum_i p_i - \sum_i q_i = \sum_{i \in X} p_i - \sum_{i \in Y} q_i \\ &= \sum_{i \in X} (p_i - v) + \sum_{i \in Y} (v - q_i). \end{aligned}$$

Let  $\mathcal{A}_{v, \mathbf{s}}$  be the basic auction that on  $\mathbf{b}'$  simulates  $\mathcal{M}(\mathbf{b}, \mathbf{s})$  to compute prices  $p_i$  for each bidder  $b'_i$  and then offers them  $p_i - v$ . It is easy to see that this is truthful, monotone (since  $\mathcal{M}$  is), and gives revenue

$$\mathcal{A}_{v, \mathbf{s}}(\mathbf{b}') = \sum_{i \in X} (p_i - v).$$

Using the bound on the revenue of any monotone basic auction (Theorem 2) we get:

$$\mathbf{E} \left[ \sum_{i \in X} (p_i - v) \right] = \mathbf{E}[\mathcal{A}_{v, \mathbf{s}}(\mathbf{b}')] \leq \mathcal{F}(\mathbf{b}').$$

Combining this with the analogous argument for  $\mathbf{s}'$  we have:

$$\mathbf{E}[\mathcal{M}(\mathbf{b}, \mathbf{s})] = \mathbf{E} \left[ \sum_{i \in X} (p_i - v) \right] + \mathbf{E} \left[ \sum_{i \in Y} (v - q_i) \right] \leq \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}').$$

□

**Theorem 3.** *For any bid vectors  $\mathbf{b}$  and  $\mathbf{s}$ , any truthful monotone double auction,  $\mathcal{M}$ , has expected profit at least  $2\mathcal{F}(\mathbf{b}, \mathbf{s})$ .*

*Proof.* Find the largest  $\ell$  such that  $b_{(\ell)} \geq s_{(\ell)}$  and choose  $v \in [s_{(\ell)}, b_{(\ell)}]$ . Now we let  $\mathbf{b}'$  and  $\mathbf{s}'$  be  $b'_i = b_i - v$  and  $s'_i = v - s_i$  for  $1 \leq i \leq n$  as in Lemma 1 giving  $\mathbf{E}[\mathcal{M}(\mathbf{b}, \mathbf{s})] \leq \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}')$  for our choice of  $v$ .

Note that  $\mathcal{F}(\mathbf{b}, \mathbf{s}) = \max_i i(b_{(i)} - s_{(i)})$ . Let  $k$  be the number of winners in  $\mathcal{F}(\mathbf{b}, \mathbf{s})$ . Note that by our choice of  $v$ , we have  $b_{(k)} \geq v$  and  $s_{(k)} \leq v$ . This gives:

$$\begin{aligned} \mathcal{F}(\mathbf{b}') &= \max_i i(b_{(i)} - v) \leq \max_i i(b_{(i)} - s_{(i)}) = \mathcal{F}(\mathbf{b}, \mathbf{s}), \text{ and} \\ \mathcal{F}(\mathbf{s}') &= \max_i i(v - s_{(i)}) \leq \max_i i(b_{(i)} - s_{(i)}) = \mathcal{F}(\mathbf{b}, \mathbf{s}). \end{aligned}$$

Thus,  $\mathbf{E}[\mathcal{M}(\mathbf{b}, \mathbf{s})] \leq \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}') \leq 2\mathcal{F}(\mathbf{b}, \mathbf{s})$ . □

## 4 Reducing Competitive Double Auctions to Competitive Basic Auctions

In this section we describe a general technique for converting any  $\beta$ -competitive basic auction into a  $2\beta$ -competitive double auction. Let  $\mathcal{A}$  be a basic auction that is  $\beta$ -competitive against  $\mathcal{F}^{(2)}(\mathbf{b})$ . We assume for the discussion here that  $b_{(n)} \geq s_{(n)}$  though it is not difficult to remove this assumption. We construct a double auction mechanism  $\mathcal{M}_{\mathcal{A}}$  that is  $2\beta$ -competitive against  $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$  as follows:

1. If  $n = 2$  run the Vickery auction, output its outcome, and halt.
2. Let  $\mathbf{b}'$  and  $\mathbf{s}'$  be  $n$ -dimensional vectors with components by  $b'_i = b_i - s_{(n)}$  and  $s'_i = b_{(n)} - s_i$ . Let  $\mathbf{b}''$  (resp.  $\mathbf{s}''$ ) be  $\mathbf{b}'$  with the smallest (resp. largest) bid deleted.
3. With probability  $1/2$  run  $\mathcal{A}(\mathbf{b}'')$ . If  $i$  wins  $\mathcal{A}$  at price  $p''_i$ , buyer  $i$  wins  $\mathcal{M}_{\mathcal{A}}$  at price  $p_i = \max(b_{(n)}, p''_i + s_{(n)})$ . Let  $k$  be the number of winners in  $\mathcal{A}(\mathbf{b}'')$ . To determine the outcome for the sellers, run the  $k$ -Vickrey auction on  $\mathbf{s}$ .



4. Otherwise (with probability  $1/2$ ) run  $\mathcal{A}(\mathbf{s}'')$ . If  $i$  wins  $\mathcal{A}$  at price  $q_i''$ , seller  $i$  wins  $\mathcal{M}_{\mathcal{A}}$  at price  $q_i = \min(s_{(n)}, b_{(n)} - q_i'')$ . As in Step 2, we run a  $k$ -Vickrey auction on the buyers to determine the outcome, where  $k$  is the number of winners in  $\mathcal{A}(\mathbf{s}'')$ .

**Theorem 4.**  $\mathcal{M}_{\mathcal{A}}$  is truthful.

We omit the proof.

**Theorem 5.** If  $\mathcal{A}$  is  $\beta$ -competitive,  $\mathcal{M}_{\mathcal{A}}$  is  $2\beta$ -competitive against  $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ .

*Proof.* If  $n = 2$ ,  $\mathcal{M}_{\mathcal{A}}$  runs Vickrey and is 2-competitive. For the rest of the proof assume  $n \geq 3$ . Let  $k \geq 2$  be the number of items sold by  $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ . Thus,

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) = k(b_{(k)} - s_{(n)}) + k(b_{(n)} - s_{(k)}) - k(b_{(n)} - s_{(n)}).$$

But by definition  $\mathcal{F}^{(2)}(\mathbf{b}') \geq k(b_{(k)} - s_{(n)})$  and likewise for  $\mathbf{s}'$ , therefore

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \leq \mathcal{F}^{(2)}(\mathbf{b}') + \mathcal{F}^{(2)}(\mathbf{s}') - k(b_{(n)} - s_{(n)}). \quad (1)$$

Note that for the buyers (and similarly for sellers):

$$\mathcal{F}^{(2)}(\mathbf{b}') \leq \mathcal{F}^{(2)}(\mathbf{b}'') + b_{(n)} - s_{(n)}. \quad (2)$$

Because  $k \geq 2$ , from Equations (1) and (2) we have

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \leq \mathcal{F}^{(2)}(\mathbf{b}'') + \mathcal{F}^{(2)}(\mathbf{s}'').$$

Note that because  $\mathcal{A}$  is  $\beta$ -competitive, the expected revenue from Step 3 and Step 4 are  $\mathcal{F}^{(2)}(\mathbf{b}'')/2\beta$  and  $\mathcal{F}^{(2)}(\mathbf{s}'')/2\beta$  respectively. Thus,

$$\mathbf{E}[\mathcal{M}_{\mathcal{A}}(\mathbf{b}, \mathbf{s})] \geq \frac{1}{2\beta}(\mathcal{F}^{(2)}(\mathbf{b}'') + \mathcal{F}^{(2)}(\mathbf{s}'')) \geq \frac{1}{2\beta}\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}).$$

□

Plugging in the 4-competitive Sampling Cost Sharing auction [6], we get a double auction with a competitive ratio of 8. Plugging in the 3.39-competitive Consensus Revenue Estimate auction [8] we get a competitive ratio of 6.78. We can do better if we customize mechanisms for the double auction problem.

## 5 The Revenue Extraction and Estimation Technique

In this section we show how the revenue extraction and revenue estimation techniques from basic auctions can be extended to the double auction problem. First we motivate the *revenue extraction* problem with a review of the Sampling Cost Sharing auction. Then we review the *consensus estimate* problem and show how it and revenue extraction combine to give a 3.75-competitive double auction.

## 5.1 Revenue Extraction

For the basic auction problem, the cost sharing mechanism of Moulin and Shenker [16] is the basis for auctions with good competitive ratios [6,8]. The cost sharing mechanism is defined as follows:

**CostShare<sub>C</sub>**: Given bids  $\mathbf{b}$ , find the largest  $k$  such that the highest  $k$  bidders can equally share the cost  $C$ . Charge each one of those  $C/k$ .

This mechanism is truthful and if  $C \leq \mathcal{F}(\mathbf{b})$  then **CostShare<sub>C</sub>** has revenue  $C$ , otherwise it has no revenue.

The SCS (Sampling Cost Sharing) auction for the basic problem is defined as follows. First partition the bidders into two sets,  $\mathbf{b}'$  and  $\mathbf{b}''$ , and compute the optimal fixed price revenues from each set,  $\mathcal{F}(\mathbf{b}')$  and  $\mathcal{F}(\mathbf{b}'')$ . Then cost share the optimal revenue for one set on the bids among the other set and vice versa (i.e., run **CostShare<sub>F(b')</sub>**( $\mathbf{b}''$ ) and **CostShare<sub>F(b'')</sub>**( $\mathbf{b}'$ )). It is easy to see that profit of the auctioneer is the minimum of the two optimal revenues. The key to the analysis is showing that the expected value of the smaller optimal revenue is at least 1/4 of the optimal revenue for  $\mathbf{b}$ .

We could attempt to follow the same mechanism framework for the double auction problem if we had an equivalent of **CostShare<sub>C</sub>** for double auctions. Unfortunately, there is no exact cost sharing analog.

**Lemma 2.** *For any value  $C$ , there is no truthful mechanism for the double auction problem that always achieves a profit of at least  $C$  when  $C$  is possible, i.e.,  $\mathcal{F}(\mathbf{b}, \mathbf{s}) \geq C$ .*

*Proof.* Suppose for a contradiction that such a mechanism  $\mathcal{M}_C$  did exist. Consider the single buyer, single seller case with  $b_1 = s_1 + C$ . Theorem 1 and  $\mathcal{M}_C$ 's truthfulness implies that the price for  $b_1$  is given by a bid-independent function  $f(s_1)$ . Since the  $C = \mathcal{F}(\mathbf{b}, \mathbf{s})$  is possible and  $\mathcal{M}_C$  is assumed to achieve at least  $C$  if it is possible,  $f(s_1)$  must be  $s_1 + C$ . Symmetrically,  $g(b_1)$  must be  $b_1 - C$ . Thus, for  $s_1$  and  $b_1$  satisfying  $b_1 = s_1 + C$  the profit of  $\mathcal{M}_C$  is  $C$  as it should be. Now consider inputs  $b'_1 = s'_1 + 2C$ . Given  $f$  and  $g$  as above, the buy price is  $p'_1 = b'_1 - C$  and the sell price is  $q'_1 = s'_1 + C = b'_1 - C$ , thus the profit is 0 giving a contradiction.  $\square$

Further, the cost sharing problem has requirements that are unnecessary for our application to profit maximizing auctions. We isolate our desired properties in the *revenue extraction problem*.

**Definition 8.** *Given a target revenue  $R$ , we want a truthful mechanism to achieve (or approximate) revenue  $R$  if the optimal profit is at least  $R$ .*

In the case of double auctions, the optimal profit above is  $\mathcal{F}(\mathbf{b}, \mathbf{s})$ . Note that unlike the cost sharing problem we do not require anything of our mechanism if  $R$  is not achievable. Furthermore, we are interested in both exact and approximate solutions.

An exact revenue extractor does not exist for the double auction problem so we define the following approximate revenue extractor that gives revenue  $\frac{k-1}{k}R$ , where  $k$  is the number of winners in  $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ . It is a natural hybrid of the  $k$ -item Vickrey [23] auction, which sells the  $k$  items to the highest  $k$  bidders at the  $(k+1)$ -st highest price, and the Moulin and Shenker cost sharing mechanism above.

**RevenueExtract<sub>R</sub>**: Given bids  $\mathbf{b}$  and  $\mathbf{s}$ , find the largest  $k$  such that  $k(b_{(k)} - s_{(k)}) \geq R$ , i.e., the  $k$  extremal buyers and sellers can generate the revenue of  $R$ . Sell to the highest  $k-1$  buyers at price  $b_{(k)}$  and buy from the lowest  $k-1$  sellers at price  $s_{(k)}$ . All other bidders including  $b_{(k)}$  and  $s_{(k)}$  are rejected.

It is easy to see that RevenueExtract<sub>R</sub> has the claimed properties.

One can convert the Sampling Cost Sharing auction for the basic problem into a double auction by using RevenueExtract<sub>R</sub> instead of CostShare<sub>C</sub>. The resulting auction is simple to describe, but its analysis is complicated by the fact that RevenueExtract<sub>R</sub> may have no revenue if there is only one item exchanged in the optimal solution. We omit its analysis because the Consensus Revenue Estimate double auction that we present next gives a better competitive ratio. None the less, the sampling cost sharing approach with revenue extraction is interesting because it appears quite general and may work in other contexts.

## 5.2 Consensus Revenue Estimate

Another application of the RevenueExtract<sub>R</sub> leads to the main result of this section: an extension of the Consensus Revenue Estimate (CORE) basic auction to the double auction problem. The resulting CORE double auction is 3.75-competitive against  $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ . (The basic CORE auction is 3.39 competitive.)

Next we describe the CORE double auction. The only differences between CORE for basic auctions [8] and CORE for double auctions are the use of RevenueExtract<sub>R</sub> instead of CostShare<sub>C</sub> and the recomputation of the optimal choice of constants  $p$  and  $c$  to take into account the fact that RevenueExtract<sub>R</sub> is an approximation.

First we describe the consensus estimate problem. For values  $r$  and  $\rho$ , function  $g$  is a  $\rho$ -consensus estimate of  $r$  if

- $g$  is a *consensus*: for any  $w$  such that  $r/\rho \leq w \leq r$ , we have  $g(w) = g(r)$ .
- $g(r)$  is a nontrivial lower bound on  $r$ , i.e.,  $0 < g(r) \leq r$ .

We define the *payoff*,  $\gamma$ , of a function  $g$  on  $r$  as

$$\gamma(r) = \begin{cases} g(r) & \text{if } g \text{ is a } \rho\text{-consensus estimate of } r \\ 0 & \text{otherwise.} \end{cases}$$

**Definition 9.** *The consensus estimate problem for  $\rho$  is to find a probability distribution,  $\mathcal{G}$ , over functions such that for all  $r$  the expected payoff,  $\mathbf{E}[\gamma(r)]$ , is big relative to  $r$ .*

The solution of [8] chooses  $\mathcal{G}$  as the following probability distribution that depends on a parameter  $c > \rho$ . Let

$$g_u(r) = r \text{ rounded down to nearest } c^{i+u} \text{ for integer } i.$$

and take  $\mathcal{G}$  as the the distribution of functions  $g_U$  for  $U$  uniform  $[0, 1]$ .

**Theorem 6.** [8] *For  $\mathcal{G}$  as defined above, for all  $r$ ,  $\mathbf{E}[\gamma(r)] = \frac{r}{\ln c}(\frac{1}{\rho} - \frac{1}{c})$ .*

It is easy to see that if  $k$ , the number of winners in  $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ , is at least three then  $\mathcal{F}^{(2)}(\mathbf{b}_{-i}, \mathbf{s})$  and  $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}_{-i})$  are in the interval  $[\frac{k-1}{k}\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}), \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})]$ .

Our CORE double auction picks  $g$  from  $\mathcal{G}$  as above and runs the bid-independent auction defined by function  $f$  (the bid-independent function for sellers is analogous):

$$f(\mathbf{b}_{-i}, \mathbf{s}) = \text{extract}_{g(\mathcal{F}^{(2)}(\mathbf{b}_{-i}, \mathbf{s}))}(\mathbf{b}_{-i}, \mathbf{s})$$

where  $\text{extract}_R$  is the bid-independent function defining the RevenueExtract $_R$  mechanism (Theorem 1 implies that  $\text{extract}_R$  exists).

Combining  $\frac{k-1}{k}$ -approximate revenue extraction with the expected payoff of consensus estimate for  $\rho = \frac{k}{k-1}$  gives the expected revenue of

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \frac{k-1}{k \ln c} \left( \frac{k-1}{k} - \frac{1}{c} \right). \quad (3)$$

Thus, we are competitive for  $k \geq 3$ .

In order to be competitive in general we must also consider the case where the number of winners in  $\mathcal{F}^{(2)}$  is  $k = 2$  (recall that  $k \geq 2$  by definition). In this case the 1-item Vickrey double auction is 2-competitive. To get an auction that is competitive for all  $k \geq 2$ , we run Vickrey with probability  $p$  and the consensus revenue estimate auction otherwise. We optimize the choice of  $p$  and  $c$  to give the CORE double auction.

**Theorem 7.** *The CORE double auction is 3.75-competitive against  $\mathcal{F}^{(2)}$ .*

*Proof.* We consider the case  $k = 2$  and  $k \geq 3$  separately.

$k = 2$ : Our expected profit is  $p\mathcal{F}^{(2)}/2$ .

$k \geq 3$ : From Vickrey we get  $p\mathcal{F}^{(2)}/k$  and from and the consensus revenue estimate we get  $(1-p)$  times the quantity in Equation (3) for a combined expected profit of:

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \left( \frac{p}{k} + \frac{(p-1)(k-1)}{k \ln c} \left( \frac{k-1}{k} - \frac{1}{c} \right) \right).$$

Our choice of  $p$  and  $c$  optimizes and balances the two cases. Numerical analysis gives  $c = 2.62$  and  $p = 0.54$  as a near-optimal choice. This choice gives a competitive ratio of 3.75.  $\square$

Note that the competitive ratio of the CORE basic auction is better than the competitive ratio of the CORE double auction (3.39 vs. 3.75). This difference is due to the fact that the former uses an exact revenue extractor and the latter uses an approximation.

## References

1. 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.
2. A. Blum, T. Sandholm, and M. Zinkevich. Online algorithms for market clearing. In *Proc. 13th Symp. on Discrete Alg.*, pages 971–980, 2002.
3. J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
4. E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.
5. S. DeVries and R. Vohra. Combinatorial Auctions: A survey. Unpublished manuscript, 2000.
6. A. Fiat, A. V. Goldberg, J. D. Hartline, and A. Karlin. Competitive generalized auctions. In *Proc. 34rd ACM Symposium on the Theory of Computing*. ACM Press, 2002. To appear.
7. Daniel Friedman and John Rust, editors. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison Wesley, 1993.
8. A. V. Goldberg and J. D. Hartline. Competitiveness via consensus. Technical Report MSR-TR-2002-73, Microsoft Research, Mountain View, CA., 2002.
9. A. V. Goldberg, J. D. Hartline, A. Karlin, and A. Wright. Competitive Auctions. Submitted to *Games and Economic Behavior*., 2001.
10. 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. Also available as InterTrust Technical Report STAR-TR-99.09.01, 1999, <http://www.star-lab.com/tr/tr-99-01.html>.
11. T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.
12. 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.
13. P. Klemperer. Auction theory: A guide to the literature. In *Journal of Economic Surveys*, pages 227–286. 13(3), 1999.
14. A. Mas-Collel, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
15. R. Preston McAfee. A dominant strategy double auction. In *Journal of Economic Theory*, volume 56, pages 434–450, 1992.
16. H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs. to appear in *Economic Theory*.
17. R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
18. N. Nisan. Algorithms for Selfish Agents: Mechanism Design for Distributed Computation. In *16th Annual Symp. on Theoretical Aspects of Comp. Sci.*, 1999.
19. 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.
20. N. Nisan and A. Ronen. Computationally feasible VCG Mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.
21. C. Papadimitriou. Algorithms, Games and the Internet. In *Proc. of 33rd Symposium Theory of Computing*, pages 749–752. ACM Press, 2001.
22. T. Sandholm and S. Suri. Market clearability. In *Proc. of the 17th International Joint Conf. on Artificial Intelligence (IJCAI)*, pages 1145–1151, 2001.
23. W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.
24. P. Wurman, W. Walsh, and M. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. In *Decision Support Systems*, volume 24, pages 17–27, 1998.