Truthful and Competitive Double Auctions

Kaustubh Deshmukh[1], Andrew V. Goldberg[2],
Jason D. Hartline[1], and Anna R. Karlin[1]

[1] Computer Science Department, University of Washington. Supported in part by NSF grant CCR-0105406. {kd,hartline,karlin}@cs.washington.edu
[2] 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 the total payments to the sellers. We aim to maximize this profit. We extend the competitive analysis framework of basic auctions [12] 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. In doing so, we generalize the consensus and revenue estimate technique from [11] to a wider class of problems.

1 Introduction

Dynamic pricing mechanisms, and specifically auctions with multiple buyers and sellers, are becoming increasing 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 (for example, see [10]). 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 [25,28].

For double auctions, the auctioneer, acting as a broker, is faced with the task of matching up a subset of the buyers with an equal-sized subset of the sellers. 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 buyer wishes to purchase exactly one item, each seller wishes to sell exactly one item, and that the items are indistinguishable, i.e., no buyer has reason to prefer one seller's item over that of another.

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 buy the item for. 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 assuming bidder utility values are independent and drawn from a known prior distribution (e.g., [7,20]). One can relax this assumption, in particular by assuming that the distribution is unknown but belongs to a certain class of distributions. Baliga and Vohra [2] discuss a number of other ways to relax this assumption as well. In this paper, following [13,9,12], we attempt to design mechanisms that maximize profit while making as few assumptions about the input as possible. As in competitive analysis of online algorithms (see e.g., [6]), 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. Note that a mechanism with a certain performance guarantee under minimal assumptions about the input distribution may have better performance guarantees under stronger assumptions, e.g., Bayesian.

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_{opt} for the buyers and a single price s_{opt} for the sellers. It then finds the largest k such that the highest k buyers each bid at least b_{opt} and the lowest k sellers each bid at most s_{opt}. It then matches these buyers and sellers up, paying all the sellers s_{opt} and charging each of the buyers b_{opt}. The profit of the auctioneer is thus $k(b_{\text{opt}} - s_{\text{opt}})$.

An interesting special case that has received attention is the *basic auction*, in which there is an auctioneer with an unlimited supply of identical items, and a set of bidders each interested in one item and each with a private utility value for the item. Previous research [13,9] has shown how to design truthful basic auctions that are competitive.

1.1 Results

This paper makes the following contributions. First 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 profit. In particular, we show that no monotone[1] double auction (even a *multi-priced* mechanism)

[1] See Section 2.2 for the definition of monotone.

can achieve a higher profit than twice the optimal *single-priced* mechanism $\mathcal{F}$ discussed above.

Next 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. Using this reduction and the best known competitive basic auction gives a double auction with a competitive ratio of 6.78. This shows that mechanism design techniques for the basic auction problem extend to the double auction problem.

Finally, we show how to apply the consensus and revenue estimate technique from [11] to the double auction directly, without using the reduction. In order to to this, we generalize the technique and make it applicable in a more general context. The generalized technique gives a 3.75-competitive double auction. We think that this technique is interesting in itself and we hope that it will find other applications.

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 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 is 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, McAffee [18] gives a truthful mechanism that approaches optimal as the number of sold items in the optimal solution grows. Note that the Vickrey-Clarke-Groves [8,14,27] 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*[2]).

Our results are closely related to the *basic auctions* for a single item available in unlimited supply, e.g., for digital goods [13,12]. As such, the approach we take in this paper closely parallels that in [12]. 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 [5] (also known as a *continuous double auction* [28]). 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

[2] Defined in Section 2.

problem was first considered in a competitive framework for profit maximization by Bar-Yossef et al. [3] and recently extended by Blum et al. [4].

More generally, there has been a great deal of recent work at the intersection of game theory, economic theory and theoretical computer science [21,24]. 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., [17], 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 [15,21,22,23]. One of the first examples of such work is the paper of Nisan and Ronen [22] 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 [26,16]).

2 Preliminaries

We consider single-round, sealed-bid double auction mechanisms in which each bidders 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 a *buyer* or a *seller*. 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 ith component of $\mathbf{b}$ (resp. $\mathbf{s}$) is b_i (resp. s_i), the bid value submitted by the ith 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.

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

- *Given the two bid vectors* $\mathbf{b} = (b_1, \ldots, b_n)$ *and* $\mathbf{s} = (s_1, \ldots, s_n)$*, 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.,* $\sum_i x_i = \sum_i y_i$.[3]

[3] 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.

  - $0 \le p_i \le b_i$ *(resp.* $s_i \le q_i$*) for all winning buyers (resp. sellers) and* $p_i = 0$ *(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., [19].*
- *If* $x_i = 1$ *buyer* i wins *(i.e., receives the item) and pays price* p_i*, otherwise we say that buyer* i loses *or is* rejected*. If* $y_i = 1$ *seller* i wins *(i.e., sells the item) and gets paid* q_i*, otherwise we say that seller* i loses *or is* rejected.
- *The* profit of the mechanism *is* $\mathcal{M}(\mathbf{b}, \mathbf{s}) = \sum_i p_i - \sum_i q_i$.

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 then 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.
- 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. (So the auctioneer does not have a strategy.)
- Bidders do not collude.

Finally, we formally define the notion of truthfulness.

Definition 2 (Truthfulness). *We say that a deterministic double auction is* truthful *if bidding truthfully, i.e.,* $b_i = u_i$ *for buyers and* $s_i = v_i$ *for sellers, is a dominant strategy for each bidder: for any value of other bidders bids, a bidder's profit (*$x_i u_i - p_i$ *for buyers and* $q_i - y_i v_i$ *for sellers) is maximized by bidding their utility.*

Definition 3. *We say that a randomized double auction is* truthful *if it can be described as a probability distribution over deterministic truthful double 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 $s_i = v_i$.

Definition 4 (Basic Auction). *[12] The* basic auction *on* n *buyers,* $\mathbf{b}$*, is a single round, sealed bid mechanism that computes an allocation* $\mathbf{x}$ *and prices* $\mathbf{p}$ *that results in the sale of up to* n *identical items.*

It is easy to see that the *basic auction* problem can be viewed as a special case of the double auction problem with all sell bids at value zero.

It is assumed that input order of the bids $\mathbf{b}$ and $\mathbf{s}$ is arbitrary. Throughout our discussion of auction problems we will use the following notation for the ith highest bidding buyer and the ith lowest bidding seller.

Definition 5. *The ith highest bidding buyer is $b_{(i)}$. The ith smallest seller bid is $s_{(i)}$.*

The most common example of a truthful basic auction is the Vickrey auction [27]. Below we give a formal definition of the Vickrey auction and extend it to the double auction problem. It is easy to see that for any fixed k, the k-Vickrey auctions are truthful. However, for any fixed k, $\mathcal{V}_k$ does not generally maximize the profit of the auctioneer.

Definition 6 (Vickrey Basic Auction). *The k-item* Vickrey basic auction *on bids* $\mathbf{b}$*,* $\mathcal{V}_k(\mathbf{b})$*, sells to the highest* k *bidders at the* $k+1$*st bid value, i.e.,* $b_{(k)}$*. Its revenue is thus*

$$\mathcal{V}_k(\mathbf{b}) = kb_{(k+1)}.$$

Definition 7 (Vickrey Double Auction[4]). *The k-item* Vickrey double auction *on bids* $\mathbf{b}$ *and* $\mathbf{s}$*,* $\mathcal{V}_k(\mathbf{b}, \mathbf{s})$*, sells to the highest* k *buyers at price* $b_{(k+1)}$ *and buys from the lowest* k *sellers at price* $s_{(k+1)}$*. Its revenue is*

$$\mathcal{V}_k(\mathbf{b}, \mathbf{s}) = k(b_{(k+1)} - s_{(k+1)}).$$

In typical examples it is important that the k-Vickrey basic auction only sells k items. It is likewise important that the k-Vickrey double auction has the same number of winning buyers as winning sellers, e.g. k of them. As specified above, the k-Vickrey auctions are not well defined if there are several bidders with identical bid values. It is simple to fix this problem by assuming that the k-Vickrey auction breaks ties arbitrarily. This tie breaking is natural and necessary for double auction problem and we will be assuming throughout this paper that bid values are distinct. This can be achieved by assuming an arbitrary (or random) total order on the bidders that respects the partial order given by their actual bid values. Thus for example, we number the bids from 1 to n, and if two bids are the same, the one with the higher number is treated as being larger. Thus, in the remainder of this paper, we will assume that $b_{(i)} > b_{(i+1)}$ and $s_{(i)} < s_{(i+1)}$.

2.1 Bid Independence

We describe a useful characterization of truthful mechanisms using the notion of *bid independence*.

Definition 8. *Let* f *and* g *be a functions from bid vectors to prices (non-negative real numbers). The* deterministic bid-independent double auction defined by f and g *is* $\mathcal{M}_{f,g}$*. For each buyer* i*,*

1. *Compute bid-independent threshold* $t_i = f(\mathbf{b}_{-i}, \mathbf{s})$*.*
   (where $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n)$*)*
2. *If* $t_i < b_i$*, set* $x_i \leftarrow 1$ *and* $p_i \leftarrow t_i$*. (Buyer* i *wins.)*

[4] This is a slight abuse of terms – The Vickrey auction for 2-sided markets will execute the trades that are economically efficient to execute, which could be less than k

3. *If* $t_i > b_i$ *set* $x_i = p_i = 0$*. (Buyer* i *is rejected.)*
4. *Otherwise, if* $t_i = b_i$ *the auction can either accept the bid at price* t_i *or reject it.*

We treat the sellers symmetrically using threshold $g(\mathbf{b}, \mathbf{s}_{-i})$ *and selling to seller* i *if the threshold is more than* s_i*.*

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

The following folklore theorem, which is a straightforward generalization of the equivalent result for basic auctions in [12], relates bid independence to truthfulness.

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

2.2 Monotonicity

We define the notion of monotone double auctions to characterize "reasonable" truthful mechanisms. The intuition underlying our notion of monotonicity is that if an auction is to achieve a large profit, the bid-independent function defining the auction should output higher prices for buyers when it sees higher buyer bid values.

Definition 9. *A randomized double auction is* monotone *if:*

- *For any pair of buyers* i *and* j *with* $b_i \leq b_j$*, we have:*

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

- *For any pair of sellers* i *and* j *with* $s_i \geq s_j$*, we have:*

$$\forall x \geq s_i,\ \mathbf{Pr}[\textit{seller } i \textit{ wins at price} \geq x] \leq \mathbf{Pr}[\textit{seller } j \textit{ wins at price} \geq x]\,.$$

To get a feel for this definition, observe that when $b_i \leq b_j$ the bids visible in the vector $\mathbf{b}_{-j}$ are the same as those visible in the 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. A similar intuition follows 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 is $\mathcal{T}(\mathbf{b}, \mathbf{s}) = \sum_i b_i - \sum_i s_i$, if all buyers bid above all sellers.

Definition 10. *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*

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

See Figure 4 for an illustration.

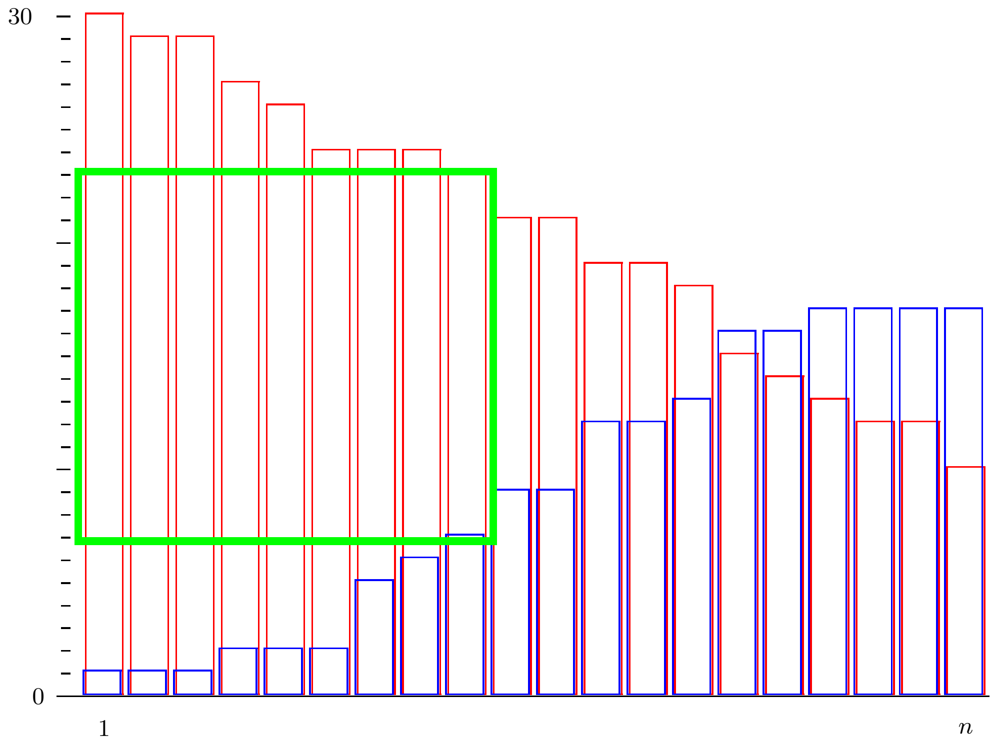

Figure 1. In this figure, the red bars represent the buyer bids sorted in order of decreasing value, the blue bars represent the seller bids sorted in order of increasing value and the area of the green rectangle represents the value of $\mathcal{F}(\mathbf{b}, \mathbf{s})$.

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 [12]) 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 [12] 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 11. *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 12. *We say that a truthful randomized double auction $\mathcal{M}$ is β-*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 β-competitive for some constant β.

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

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 buyer and seller 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:

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

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

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

□

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

Proof. Find the largest ℓ such that $b_{(\ell)} \ge s_{(\ell)}$ 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 $1 \le i \le n$. For our choice of v, Lemma 1 gives $\mathbf{E}[\mathcal{M}(\mathbf{b},\mathbf{s})] \le \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}')$.

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 $b_{(k)} \ge v$ and $s_{(k)} \le v$. This gives:

$$\mathcal{F}(\mathbf{b}') = \max_i i(b_{(i)} - v) \le \max_i i(b_{(i)} - s_{(i)}) = \mathcal{F}(\mathbf{b},\mathbf{s}), \text{ and}$$
$$\mathcal{F}(\mathbf{s}') = \max_i i(v - s_{(i)}) \le \max_i i(b_{(i)} - s_{(i)}) = \mathcal{F}(\mathbf{b},\mathbf{s}).$$

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

4 Reducing Competitive Double Auctions to Competitive Basic Auctions

In previous work, a number of different constant competitive mechanisms have been described for the basic auction problem. At first glance, one might imagine that one could simply run a competitive basic auction on the buyers (with auctioneer as seller) and a competitive basic auction on the sellers (with auctioneer as buyer) and combine the results. The problem is that the number of items sold in the auction on the buyers might be different from the number of items bought in the auction on the sellers. Thus a mechanism is needed for coordinating these outcomes. In this section, we present a technique for doing this. Specifically, we describe a general procedure for converting any β-competitive basic auction into a 2β-competitive double auction.

Definition 13 **($\mathcal{M}_{\mathcal{A}}$).** *Given,*

- *Basic auction,* $\mathcal{A}$.
- *Input* $\mathbf{b}$ *and* $\mathbf{s}$.

– *ℓ, the largest value such that $b_{(\ell)} \geq s_{(\ell)}$.*

The double auction, $\mathcal{M}_{\mathcal{A}}$, does as follows:

Case 1 ($\ell = 1$): *Output the empty allocation.*

Case 2 ($\ell = 2$): *Simulate the 1-item Vickrey double auction, $\mathcal{V}_1(\mathbf{b}, \mathbf{s})$, and output its outcome.*

Case 3 ($\ell \geq 3$): *Let $\mathbf{b}'$ and $\mathbf{s}'$ be n-dimensional vectors with components by $b'_i = b_i - s_{(\ell)}$ and $s'_i = b_{(\ell)} - s_i$. Let $\mathbf{b}''$ (resp. $\mathbf{s}''$) be the $(\ell - 1)$-dimensional vector consisting of the largest $\ell - 1$ bids in $\mathbf{b}'$ (resp. $\mathbf{s}'$).*

1. *With probability $1/2$ simulate $\mathcal{A}(\mathbf{b}'')$. If buyer i wins the simulation of $\mathcal{A}$ at price p''_i, then buyer i wins $\mathcal{M}_{\mathcal{A}}$ at price $p_i = \max(b_{(\ell)}, p''_i + s_{(\ell)})$. All other buyers lose. Let $k < \ell$ be the number of winners in $\mathcal{A}(\mathbf{b}'')$. To determine the outcome for the sellers, run the k-Vickrey basic auction auction on $\mathbf{s}$.*
2. *Otherwise (with probability $1/2$) simulate $\mathcal{A}(\mathbf{s}'')$. If seller i wins the simulation of $\mathcal{A}$ at price q''_i then seller i wins $\mathcal{M}_{\mathcal{A}}$ at price $q_i = \min(s_{(\ell)}, b_{(\ell)} - q''_i)$. As in Step 1, we run the k-Vickrey basic auction on the buyers to determine the outcome for buyers, where $k < \ell$ is the number of winners in $\mathcal{A}(\mathbf{s}'')$.*

We illustrate the auction pictorially in the following figures.

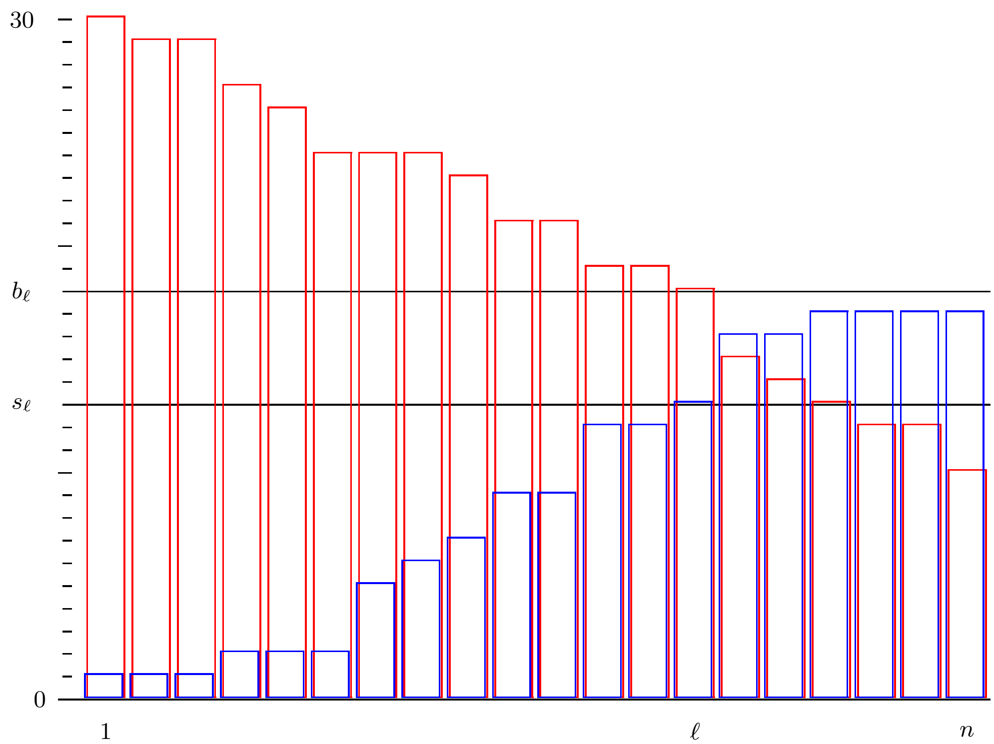

Figure 2. In this and the remaining figures, the red bars represent the buyer bids sorted in order of decreasing value, the blue bars represent the seller bids sorted in order of increasing value. The values $b_{(\ell)}$ and $s_{(\ell)}$ are shown.

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

Proof. One way to prove this is to give a bid-independent implementation. As this is rather tedious, we present a direct argument here.

We show that $\mathcal{M}_\mathcal{A}$ is truthful for the buyers. The result for sellers is symmetric. First note all buyers that win the auction pay at least $b_{(\ell)}$ and that the buyer with the ℓth highest bid and all buyers with lower bids always lose the auction. In the case that $\ell \leq 2$ this is obvious. For $\ell \geq 3$ we have:

- In Step 1, since $b_{(\ell)}$ is excluded from $\mathbf{b}''$, $b_{(\ell)}$ loses. In this case by definition, all winners pay at least $b_{(\ell)}$.
- In Step 2, since $k < \ell$ the k-Vickrey auction on buyers rejects $b_{(\ell)}$ (and winners must pay at least $b_{(\ell)}$).

We now argue that $\mathcal{M}_\mathcal{A}$ is truthful for the buyer with the ℓth highest bid and all lower bidding buyers. Suppose buyer i is one of these lower bidding buyers. Hold all other bid values fixed. To win the auction, buyer i would have to bid higher than $b_{(\ell-1)}$. In this case, all winners would pay at least $b_{(\ell-1)}$ which is more than this buyer's utility value, b_i. Such a sale price would give buyer i a negative profit, hence buyer i prefers losing the auction.

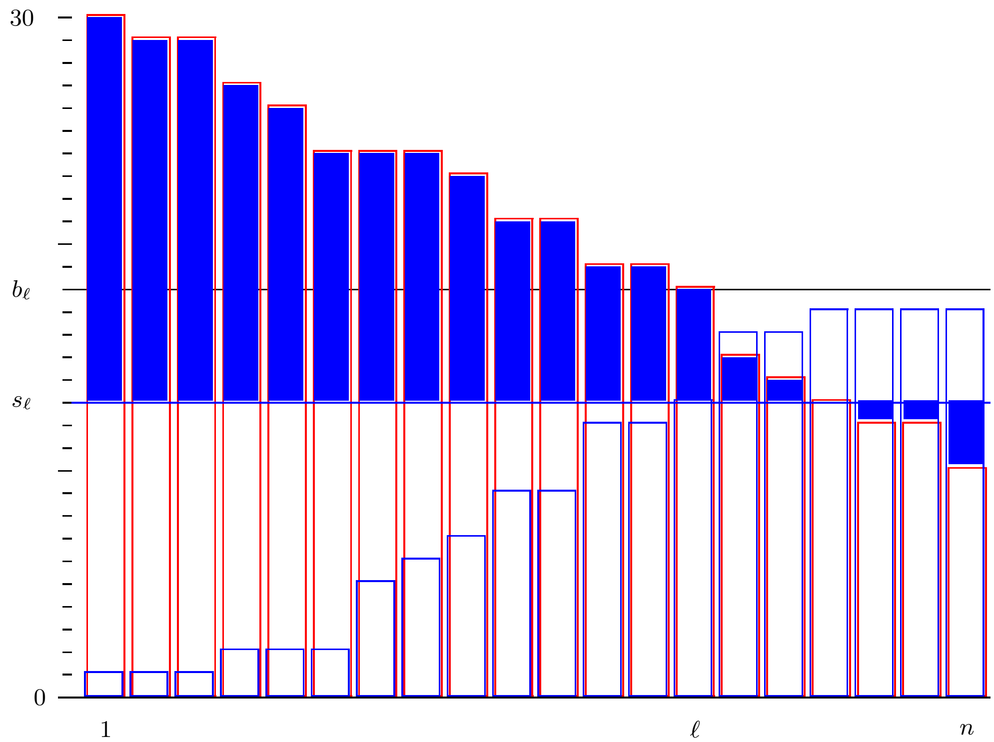

Figure 3. The solid blue bars represent the bids $b'_i = b_i - s_{(\ell)}$ sorted in order of decreasing value. In step 1 of case 3, which occurs with probability 1/2, we run the auction $\mathcal{A}$ on the first $\ell - 1$ of these (all those with $b'_i > 0$).

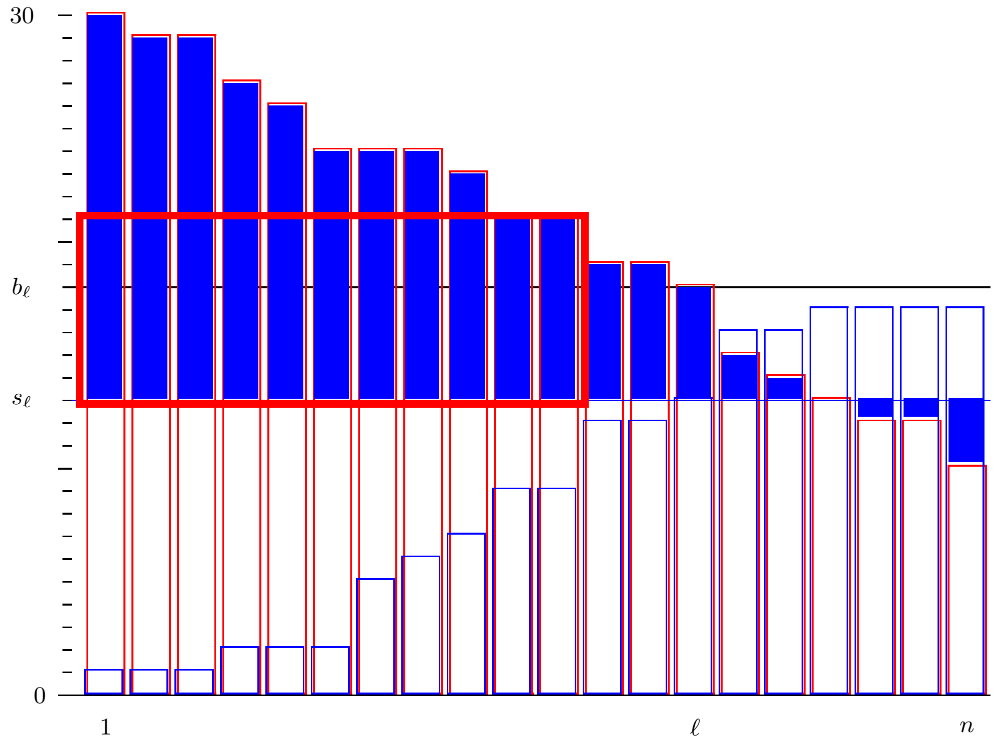

Figure 4. The solid blue bars represent the bids $b'_i = b_i - s_{(\ell)}$ sorted in order of decreasing value. In step 1 of case 3, we run the auction $\mathcal{A}$ on the first $\ell - 1$ of these (all those with $b'_i > 0$). The area of the red rectangle represents the optimal fixed price revenue for this set of bids. The auction $\mathcal{A}$ is guaranteed to return a profit from these bidders which is at least a constant fraction of this area.

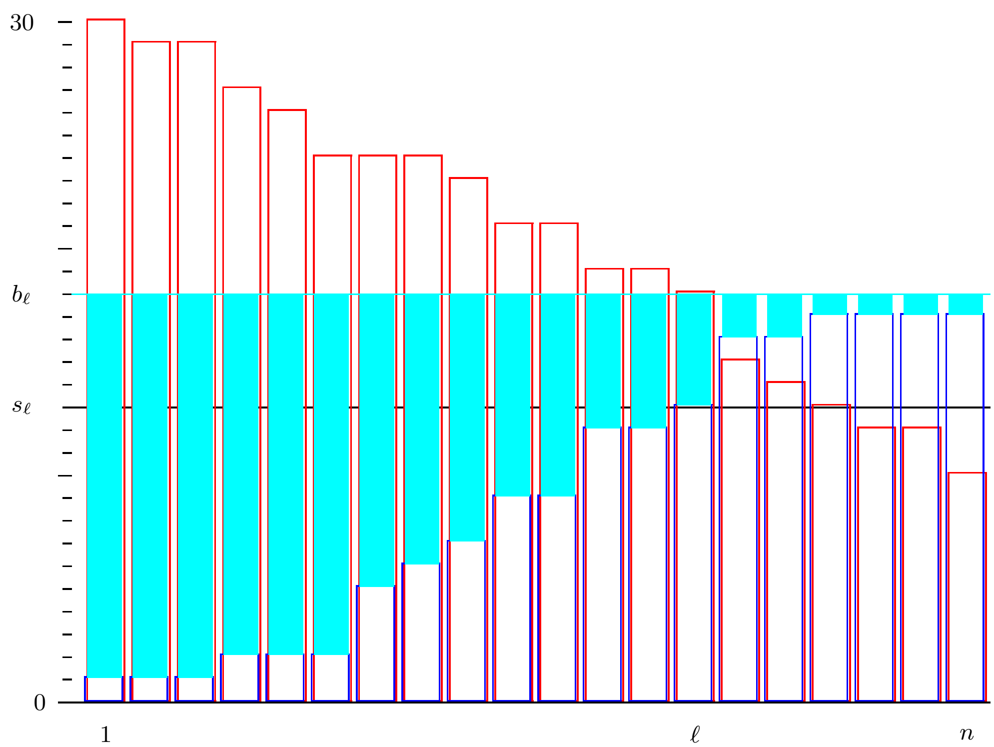

Figure 5. The solid cyan bars represent the bids $s'_i = b_{(\ell)} - s_i$ sorted in order of decreasing value. In step 2 of case 3, we run the auction $\mathcal{A}$ on the first $\ell - 1$ of these (all those with $b'_i > 0$.)

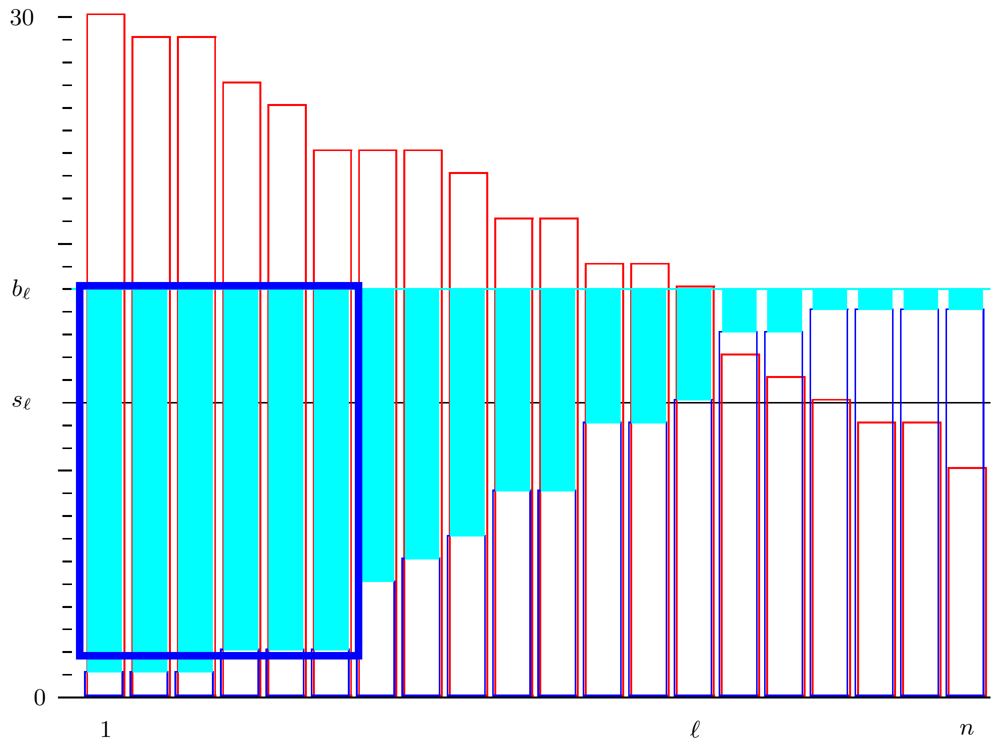

Figure 6. The solid cyan bars represent the bids $s'_i = b_{(\ell)} - s_i$ sorted in order of decreasing value. In step 2 of case 3, we run the auction $\mathcal{A}$ on the first $\ell - 1$ of these. The area of the blue rectangle represents the optimal fixed price revenue for this set of bids. The auction $\mathcal{A}$ is guaranteed to return a profit from these bidders which is at least a constant fraction of this area. Note that if a selling bidder wins the basic auction simulation at a price of p, the price paid to that seller is $\min(s_{(\ell)}, b_\ell - p)$. Since the corresponding winning buyer (determined by running a k-Vickrey auction pays at least b_ℓ, the profit associated with this selling bidder is at least p.

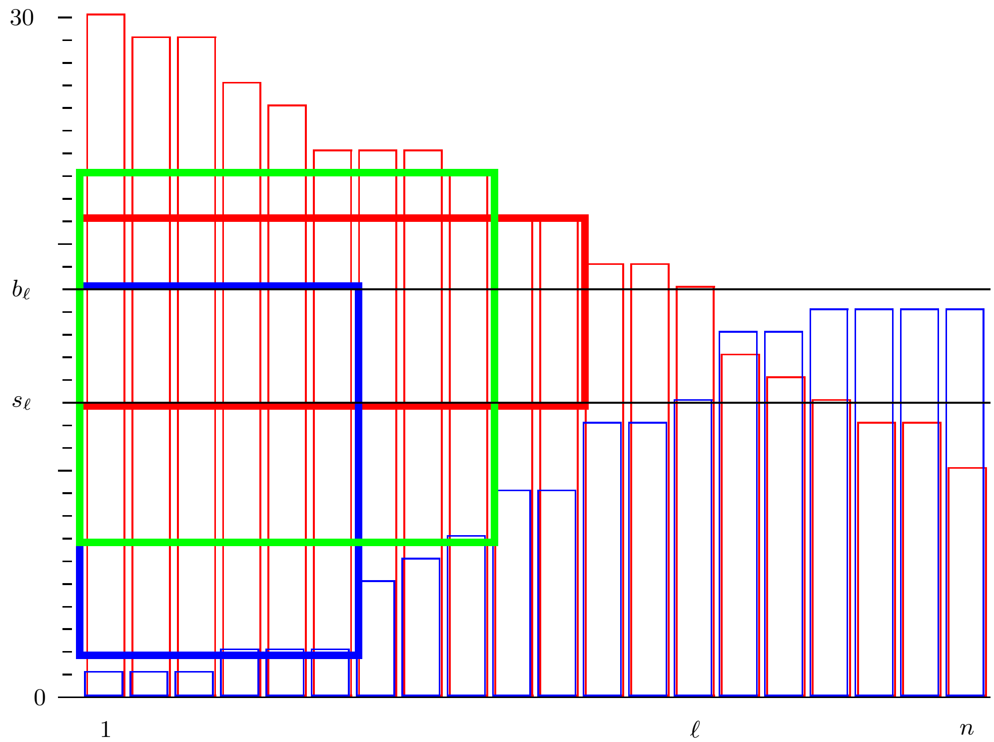

Figure 7. In this figure, the area of the red rectangle is the optimal fixed price revenue from the buyer bids $\mathbf{b}'$ and the area of the blue rectangle is the optimal fixed price revenue from the seller bids $\mathbf{s}'$. The green rectangle is the optimal fixed price revenue in the double auction. The area of this rectangle is at most the area of the red rectangle plus the area of the blue rectangle.

Now we show that the mechanism is truthful for the remaining $\ell - 1$ high bidding buyers. First, none of these bidders can change the value of $b_{(\ell)}$ or ℓ without lowering their bid value below $b_{(\ell)}$ which would cause them to lose the auction.

We now argue (the simpler case) that Step 2 is truthful for the $\ell - 1$ high bidding buyers. In this step the truthful k-Vickrey auction is run on these buyers. Since k-Vickrey is truthful and because the k is determined from $\mathbf{s}$, ℓ, and $b_{(\ell)}$ which we have shown to be unchangeable by any winning buyers, Step 2 is truthful.

The truthfulness of Step 1 is similar. Suppose basic auction $\mathcal{A}$ is defined by bid-independent function f'. Then we can define the bid-independent function for this case as $f(\mathbf{b}_{-i}, \mathbf{s}) = \max(b_{(\ell)}, f'(\mathbf{b}_{-i} - s_{(\ell)}) + s_{(\ell)})$. Given that $b_{(\ell)}$, ℓ, and thus $s_{(\ell)}$ have been shown to be unchangeable by any winning buyers, this shows that Step 1 is truthful. □

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

Proof. If $\ell = 1$, $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \leq 0$ so the null allocation is competitive. If $\ell = 2$, $\mathcal{M}_{\mathcal{A}}$ runs the 1-item Vickrey double auction which is 2-competitive when $\ell = 2$. For the rest of the proof assume $\ell \geq 3$. Let $k \in [2, \ell]$ be the number of items sold by $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$. Thus,

$$\begin{aligned} \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) &= k(b_{(k)} - s_{(k)}) \\ &= k(b_{(k)} - s_{(\ell)}) + k(b_{(\ell)} - s_{(k)}) - k(b_{(\ell)} - s_{(\ell)}). \end{aligned}$$

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

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

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

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

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 β-competitive, the expected revenues from Step 3 and Step 4 are at least $\mathcal{F}^{(2)}(\mathbf{b}'')/2\beta$ and $\mathcal{F}^{(2)}(\mathbf{s}'')/2\beta$ respectively. Thus,

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

□

Plugging in the 4-competitive Sampling Cost Sharing (CORE) basic auction [9], we get a double auction with a competitive ratio of 8. Plugging in the 3.39-competitive Consensus Revenue Estimate basic auction [11] we get a competitive ratio of 6.78. As we show below, we can do better if we customize mechanisms for the double auction problem.

5 Generalized CORE Mechanism

As we have seen in the previous section, the CORE mechanism of [11] can be used in the general reduction from the basic auction problem to give a 6.78-competitive auction for the double auction problem. In this section, we show how the ideas of the CORE auction can be generalized to get a double auction mechanism with a competitive ratio of 3.75.

We begin with an intuitive explanation of the essential ideas of CORE as applied in the basic auction setting. The mechanism relies on the fact that, given any target revenue R, there exists a truthful mechanism parameterized by R that is guaranteed to extract a profit of R from the bidders as long as $R \leq \mathcal{F}(\mathbf{b})$. This observation suggests the following idea for an auction: First, try to truthfully estimate a value R that is close to, but no greater than $\mathcal{F}(\mathbf{b})$. Then run the truthful mechanism that extracts the profit of R from the bidders. This is the essential idea of CORE.

A bit more formally, CORE consists of two parts: a *consensus revenue estimator* and a *profit extractor*.

1. Get a bid-independent *consensus estimate* of $\mathcal{F}(\mathbf{b})$.
   This consists of selecting a bid-independent function $r(\cdot)$ (according to some predefined probability distribution over such functions) so that:
   - *The outcome is a consensus::* This means that with high probability, $r(\mathbf{b}_{-i}) = R$ for all i, i.e., even though the computation performed for each bidder is on a different subset of the bids (specifically, all bids but his own), the computations for the different bidders all compute the same value R.
   - *The outcome is a good estimate:* This means that R is close to but less than $\mathcal{F}(\mathbf{b})$.
2. Truthfully extract the profit R. (Below, we will refer to the bid-independent function implementing this profit extraction as $\text{pe}_R(\cdot)$.)

To apply this approach to the double auction setting, we would like a consensus revenue estimator that gives us a good estimate of $\mathcal{F}(\mathbf{b}, \mathbf{s})$, and a profit extractor, i.e., a truthful mechanism that extracts a revenue of R whenever $R \leq \mathcal{F}(\mathbf{b}, \mathbf{s})$.

There are two issues that prevent the generalization to double auctions from being immediate. The first is that there is no profit exact or even approximate profit extractor which extracts a revenue of R whenever $R \leq \mathcal{F}(\mathbf{b}, \mathbf{s})$. The second is that for the basic auction, there was no problem in the unlikely event that the first step fails and consensus is not achieved. In contrast, in the double auction setting, if consensus is not achieved, it is not automatically guaranteed, for example, that the same number of buyers and sellers win the auction. Thus, we need to show that the outcome of the auction is never infeasible.

In the rest of this section, we formalize the ideas that have to this point been described informally. We present an abstract version of CORE that has the potential to be useful for a wide variety of private-value optimization problems.

We begin by describing the class of mechanism design problems that we consider and then proceed to give a general version of the CORE mechanism. Once we have presented the general technique, we apply it to the double auction problem.

5.1 Single Parameter Agent Allocation Problems

We consider mechanism design problems with the following characteristics:

- There are n agents.
- The feasible outcomes of the mechanism are a subset $\mathcal{X}$ of the possible allocations $\mathbf{x} = (x_1, \ldots, x_n)$, where x_i is 1 if agent i receives the allocation, and 0 otherwise.
- Each agent has a utility value, where agent i's utility is u_i.

The mechanism takes as input a vector of bids, one from each agent, representing that agents' utility value, and produces as output an allocation $\mathbf{x} \in \mathcal{X}$ and a payment vector $\mathbf{p} \in \mathbb{R}^n$, representing the payment from the bidder to the auctioneer. We require that $p_i = 0$ if $x_i = 0$, and $p_i \leq u_i$ otherwise.

We use the same private value model as discussed before in the context of double auctions, so, for example, bidder i bids so as to maximize his profit, which is $u_i x_i - p_i$.

The auctioneer's profit is $\sum_i p_i$, and our goal is to design truthful mechanisms which maximize auctioneer profit.

This model captures both the basic auction and the double auction. The double auction is captured by (a) representing the utility value of the selling bidders as negative numbers (the negation of which is the minimum that bidder is willing to take for the item) and (b) letting $\mathcal{X}$ be the set of allocations in which an equal number of the buying bidders and selling bidders receive an allocation.

A more complex example captured by this model is a combinatorial auction in which each bidder is single-minded. In this case, there is a set of n different items to be auctioned off, each bidder is interested in only one subset of the available items, and has a utility value for that subset, and the feasible allocations are partitions of the items among the bidders, where bidder i has $x_i = 1$ if and only if he receives a superset of his desired set of items.

Since all of these problems are allocation problems with profit maximization as the objective, we simply use the term auction to describe any of them.

It is well known that for this type of auction, in which each bidder is defined by a single utility value (this case has been referred to as the *single parameter agents* case), Theorem 1 holds.

Theorem 6. *[1] An auction with single-parameter agents is truthful if and only if it is bid-independent.*

5.2 Profit Extraction

We are now ready to present the generalized CORE mechanism for any auction problem with single parameter agents. We begin by discussing the profit extraction step.

To define a profit extractor we need to define a profit metric. Let $\mathbf{I}$ represent the input to the auction problem and the *masked input* $\mathbf{I}_{-i}$ be the input with the value corresponding to the i-th bidder removed. For example $\mathbf{I} = \mathbf{b}$ for the basic auction problem and $\mathbf{I} = (\mathbf{b}, \mathbf{s})$ for the double auction problem.

Definition 14 (Profit Metric). *A* profit metric, $\mathcal{O}$, *is a function from inputs* $\mathbf{I}$ *to nonnegative reals (to be interpreted as an amount of profit achievable on that input as measured by the metric of choice).*

We use a profit metric as a benchmark for auction performance. As an example, $\mathcal{F}(\mathbf{b}, \mathbf{s})$ or $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ are two of the possible profit metrics for the double auction problem and $\mathcal{F}(\mathbf{b})$ and $\mathcal{F}^{(2)}(\mathbf{b})$ are two of the possible profit metrics for the basic auction problem.

For any $\alpha \geq 1$, we use the following definition.

Definition 15 (α-Profit Extractor). *A* α-profit extractor *for profit metric* $\mathcal{O}$ *is a truthful mechanism* $\mathcal{O}$-*Extract*$_R$ *parameterized by a target profit* R*, that has expected profit of at least* R/α *on all inputs,* $\mathbf{I}$*, such that* $\mathcal{O}(\mathbf{I}) \geq R$*.*

We refer to a 1-profit extractor as *exact*.

Intuitively, if the optimal profit on a given input $\mathbf{I}$ is small, e.g., much smaller than our target R, we do not require the profit extractor $\mathcal{O}$-Extract$_R$ to work well or, for that matter, to work at all. However, if the optimal profit is substantial, greater than our target R, we require $\mathcal{O}$-Extract$_R$ to extract a constant fraction of our target.

We illustrate this concept by discussing a known profit extractor for basic auctions, and presenting a new approximate profit extractor for double auctions.

Profit Extraction for Basic Auctions For the basic auction problem, a specialization of the general *cost sharing* mechanism of Moulin and Shenker [19] is a 1-profit extractor for metric $\mathcal{F}$. It is also the basis for the only basic auctions known to have good competitive ratios [9,11]. The cost sharing mechanism for the basic auction problem is defined as follows:

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

This mechanism is truthful and, if $R \leq \mathcal{F}(\mathbf{b})$, then CostShare$_R$ has revenue R. Otherwise it has no revenue.

Profit Extraction for Double Auctions For the double auction problem, we present a simple approximate profit extractor for metric $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$.

Definition 16. *The double auction profit extractor,* $\mathcal{F}^{(2)}$-*Extract*$_R$*, on input* $\mathbf{b}$ *and* $\mathbf{s}$ *computes the largest* k *such that* $k(b_{(k)} - s_{(k)}) \geq R$*, buys from the lowest* $k-1$ *bidding sellers at price* $s_{(k)}$*, and sells to the top* $k-1$ *bidding buyers at price* $b_{(k)}$*. All other buyers and sellers (including the* k*th) are rejected.*

Note the resemblance of this mechanism to the optimal Bayesian mechanism. The latter, knowing buyer and seller utility distributions, computes optimal prices x and y, and for pairs $b_{(i)}, s_{(i)}$ such that both $b_{(i)} \geq x$ and $s_{(i)} \leq y$, buys from $s_{(i)}$ at price y and sells to $b_{(i)}$ for price x. Our auction chooses the prices x and y adaptively, by computing R and k from the input.

We first show that this profit extractor is truthful.

Lemma 2. *$\mathcal{F}^{(2)}$-Extract$_R$ is truthful.*

Proof. We show that the auction is bid-independent for buyers. The case of sellers is analogous. Define the bid-independent function $f(\mathbf{b}_{-i}, \mathbf{s})$ for buyers as follows for input $\mathbf{b}_{-i}$ and $\mathbf{s}$:

- Compute $\mathbf{b}^{(i)}$ from $\mathbf{b}$ by replacing b_i by ∞.
- Simulate $\mathcal{F}^{(2)}$-Extract$_R$ on $\mathbf{b}^{(i)}$ and $\mathbf{s}$ to compute a buyer allocation and prices.
- If buyer i is not allocated the item in the simulation, then output price ∞.
- Otherwise output p_i, the price buyer i payed in the simulation.

We now show that this bid-independent mechanism, $\mathcal{M}_f$, is identical to $\mathcal{F}^{(2)}$-Extract$_R$.

Suppose buyer i bidding b_i is allocated an item at price p_i by $\mathcal{F}^{(2)}$-Extract$_R(\mathbf{b}, \mathbf{s})$. This occurs because the largest k such that $k(b_{(k)} - s_{(k)}) \geq R$ satisfies $b_{(k)} < b_i$. It is easy to see that the value of k does not change if we increase b_i. Thus $\mathcal{F}^{(2)}$-Extract$_R$ sells to b_i at price $b_{(k)}$.

Suppose now that buyer i bidding b_i is not allocated an item by $\mathcal{F}^{(2)}$-Extract$_R(\mathbf{b}, \mathbf{s})$. This means that for all k such that $b_{(k)} < b_i$, $k(b_{(k)}, s_{(k)}) < R$. This remains true if we increase b_i. Thus the simulation either finds a sale price that that is more than b_i or it does not sell any items. In either case buyer i is rejected by $\mathcal{M}_f$. □

The following theorem bounds the performance of this profit extractor.

Theorem 7. *On bids $\mathbf{b}$ and $\mathbf{s}$ such that $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ exchanges k items, $\mathcal{F}^{(2)}$-Extract$_R$ is a $\frac{k}{k-1}$-profit extractor for $\mathcal{F}^{(2)}$.*

Proof. Suppose that $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \geq R$. Let k be the number of items exchanged by $\mathcal{F}^{(2)}$ and let k' be the number of items exchanged by $\mathcal{F}^{(2)}$-Extract$_R$. Note that $\mathcal{F}^{(2)}$-Extract$_R$ finds the largest k' such that $k'(b_{(k')} - s_{(k')}) \geq R$. By the definition of k, $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) = k(b_{(k)} - s_{(k)}) \geq R$. Thus, $k' \geq k$. The revenue of $\mathcal{F}^{(2)}$-Extract$_R$ is

$$\begin{aligned}\mathcal{F}^{(2)}\text{-Extract}_R(\mathbf{b}, \mathbf{s}) &= (k'-1)(b_{(k')} - s_{(k')}) \\ &\geq \tfrac{k'-1}{k'} R \geq \tfrac{k-1}{k} R.\end{aligned}$$

□

Although in general profit extractors can be randomized, the applications in this paper use deterministic extractors. From now on, we assume a deterministic profit extractor $\mathcal{O}$-Extract$_R$. The characterization of truthful mechanisms as those that have bid-independent implementations (Theorem 1) implies that the truthful mechanism $\mathcal{O}$-Extract$_R$ is equivalent to an auction defined by a bid-independent function $\text{pe}_R(\cdot)$.

5.3 Consensus Revenue Estimation

We turn next to the problem of truthfully computing a value R that will be input to the profit extractor. This is done via *consensus revenue estimation.*

Fix a profit metric $\mathcal{O}$ on inputs.$\mathbf{I}$. Consider a function r from masked inputs into reals. We say that r is a *consensus estimate* if there is a value $R \leq \mathcal{O}(\mathbf{I})$ such that for all bidders i, $r(\mathbf{I}_{-i}) = R$.

Given r and pe_R, consider the auction, $\mathcal{M}$, defined by bid-independent function

$$f(\mathbf{I}_{-i}) = \text{pe}_{r(\mathbf{I}_{-i})}(\mathbf{I}_{-i}).$$

If r is a consensus estimate, then $\mathcal{M}$ is identically $\mathcal{O}$-Extract$_R$ and its profit if at least R/α.

Our consensus revenue estimation technique is based on the observation that no single bidder should have a significant effect on the achievable profit. This property holds in most cases of interest. In exceptional cases, either no truthful mechanism is competitive or we use an alternative exception-handling technique. This motivates the following definition:

Definition 17 (Sensitivity). *For $\rho > 1$, an input, $\mathbf{I}$, is ρ-insensitive for metric $\mathcal{O}$ if for all bidders, i, we have:*

$$\mathcal{O}(\mathbf{I})/\rho \leq \mathcal{O}(\mathbf{I}_{-i}) \leq \mathcal{O}(\mathbf{I}).$$

Consider a ρ-insensitive input $\mathbf{I}$ and metric $\mathcal{O}$. Following [11], where it was shown that there is no deterministic consensus estimate function, we choose the function $r(\cdot)$ according to the following probability distribution:

Definition 18 (Consensus Estimate Function Distribution). *Parameterized by $c > 1$ and the metric $\mathcal{O}$ we define the random function $r(\cdot)$ as follows:*

- *Let U be a uniform random variable from $[0, 1]$.*
- *Define $r(\mathbf{I})$ as $\mathcal{O}(\mathbf{I})$ rounded down to the nearest c^{i+U} for integer i.*

Notice that on ρ-insensitive inputs, for all i,

$$\mathcal{O}(\mathbf{I})/c\rho < r(\mathbf{I}_{-i}) \leq \mathcal{O}(\mathbf{I}),$$

and thus $r(\cdot)$ is a fairly good revenue estimator (to within a factor $c\rho$). The problem is that we have no means of coping with the situation where r fails to produce a consensus. In this latter situation, we cannot count on extracting any profit whatsover. The following lemma bounds the expected revenue of this technique, i.e., the expectation of the consensus estimate value, assuming that the value is zero if r fails to compute a consensus estimate. For constant ρ and c, this expected value is a constant fraction of $\mathcal{O}(\mathbf{I})$.

Lemma 3. *[11] For ρ-insensitive inputs $\mathbf{I}$, let random variable X be R if $r(\mathbf{I}_{-i}) = R$ for all i and zero otherwise. Choosing $r(\cdot)$ as defined by the distribution , and assuming that $c > \rho$, we have:*

$$\mathbf{E}[X] \geq \frac{\mathcal{O}(\mathbf{I})}{\ln c}\left(\frac{1}{\rho} - \frac{1}{c}\right).$$

5.4 Putting it together

We are now ready to present the general definition of CORE.

Definition 19 (CORE$_c$). *For constant c, metric $\mathcal{O}$, and profit extractor $\mathcal{O}$-Extract$_R$ for $\mathcal{O}$, defined bid-independently by* pe_R*, the mechanism CORE$_c$ is defined bid-independently by*

$$f(\mathbf{I}_{-i}) = \mathrm{pe}_{r(\mathbf{I}_{-i})}(\mathbf{I}_{-i})$$

for r sampled according to the distribution given in Definition 3.

Before we can claim that CORE performs well, we must take care of one other potential pitfall.

Definition 20 (Safety). *A profit extractor $\mathcal{O}$-Extract$_R$ is* safe *if it is impossible for CORE$_c$ to output an infeasible outcome (an allocation that is not in $\mathcal{X}$) or an outcome in which the auctioneer's profit is negative.*

Of course, by construction, CORE$_c$ can never output infeasible outcomes or ones with negative profit when there is consensus. To prove that a profit extractor is safe, one must show that no infeasible or bad situations can arise in the event that there is no consensus. For example, for double auctions, we must ensure that when there is no consensus the number of winning buyers and sellers is the same.

The performance of CORE follows from the above lemmas and definitions and is summarized in the following theorem.

Theorem 8. *Given*

- *profit metric $\mathcal{O}$,*
- *safe α approximate profit extractor $\mathcal{O}$-Extract$_R$ for $\mathcal{O}$ implemented by bid-independent function* pe_R*, and*
- *and input $\mathbf{I}$ that is ρ-sensitive for metric $\mathcal{O}$,*
- *$c > \rho$;*

the $CORE_c$ mechanism has an expected revenue and competitive ratio, respectively, of:

$$\mathbf{E}[CORE_c(\mathbf{I})] \geq \frac{\mathcal{O}(\mathbf{I})}{\alpha \ln c}\left(\frac{1}{\rho} - \frac{1}{c}\right), \qquad \frac{\mathcal{O}(\mathbf{I})}{\mathbf{E}[CORE_c(\mathbf{I})]} \geq \alpha \ln c \left(\frac{1}{\rho} - \frac{1}{c}\right)^{-1}.$$

Given ρ, the value of c can be chosen to give the optimal competitive ratio when restricted to ρ-sensitive inputs.

6 Application to Double Auctions

We now specialize the results of the previous section to the double auction setting.

For double auctions, we have already presented an approximate profit extractor for profit metric $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$. Thus, we have only two loose ends to tie up:

- We must show that $\mathcal{F}^{(2)}$-Extract$_R$ is safe.
- We must deal with the exceptional case that the input is not ρ-insensitive for some constant ρ.

6.1 Safety

Lemma 4. *$\mathcal{F}^{(2)}$-Extract$_R$ is safe.*

Proof. We must show that CORE$_c$ with $\mathcal{F}^{(2)}$-Extract$_R$ does not effect infeasible outcomes, i.e., outcomes where the number of winning buyers is not equal to the number of winning sellers. Let k be the number of items sold by $\mathcal{F}^{(2)}$ on input $\mathbf{b}$ and $\mathbf{s}$ and let $F = \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ be its profit.

1. For buyer i not allocated items by $\mathcal{F}^{(2)}$ we have $\mathcal{F}^{(2)}(\mathbf{b}_{-i}, \mathbf{s}) = F$. Likewise for sellers.
2. $\mathcal{F}^{(2)}$-Extract$_R$ for $R \leq F$ has at least k items exchanged.

Note that the top k buyers and sellers always exchange items regardless of consensus. Further, by 1 above, the bottom $n - k$ buyers and sellers always have consensus. Thus the same number of these additional buyers and sellers exchange items.

6.2 Dealing with non-sensitive inputs

As discussed above, the CORE approach only works on $\mathbf{b}$ and $\mathbf{s}$ that are ρ-sensitive for $\mathcal{F}^{(2)}$. Our goal, however, is an auction that is constant-competitive with $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ for any $\mathbf{b}$ and $\mathbf{s}$. To deal with this, we make the following observations about $\mathcal{F}^{(2)}$. Let k be the number of winning buyer-seller pairs in $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$.

– If $k \geq 3$ then $\mathbf{b}$ and $\mathbf{s}$ are $\frac{k}{k-1}$-insensitive:

$$\frac{k-1}{k}\mathcal{F}^{(2)}(\mathbf{b},\mathbf{s}) \leq \mathcal{F}^{(2)}(\mathbf{b}_{-i},\mathbf{s}), \mathcal{F}^{(2)}(\mathbf{b},\mathbf{s}_{-i}) \leq \mathcal{F}^{(2)}(\mathbf{b},\mathbf{s}).$$

In this case, for $c > 3/2$, CORE_c with profit extractor $\mathcal{F}^{(2)}$-Extract_R is constant-competitive:

$$\mathbf{E}[\text{CORE}_c(\mathbf{b},\mathbf{s})] = \frac{\mathcal{F}^{(2)}(\mathbf{b},\mathbf{s})}{\ln c}\left(\frac{k-1}{k}\right)\left(\frac{k-1}{k} - \frac{1}{c}\right). \tag{3}$$

– For $k = 2$, $\mathcal{V}_1$ is 2-competitive with $\mathcal{F}^{(2)}$.

Thus, it is possible to take a convex combination of the CORE_c auction with $\mathcal{V}_1$ to get a constant-competitive auction in worst-case.

Definition 21 ($\text{CORE}_{c,p}$). *The double auction $\text{CORE}_{c,p}$ parameterized by $c > 3/2$ and $p \in (0,1)$ with probability $1-p$ runs CORE_c with profit extractor $\mathcal{F}^{(2)}$-Extract_R and with probability p runs $\mathcal{V}_1$.*

Theorem 9. *With a near optimal choice of $c = 2.62$ and $p = 0.54$, the $\text{CORE}_{c,p}$ double auction with profit extractor $\mathcal{F}^{(2)}$-Extract_R is 3.75-competitive against $\mathcal{F}^{(2)}$.*

Proof. Let k be the number of items sold by $\mathcal{F}^{(2)}(\mathbf{b},\mathbf{s})$. We consider the case $k = 2$ and $k \geq 3$ separately.

Case 1 ($k = 2$): Our expected profit is $p\mathcal{F}^{(2)}/2$.
Case 2 ($k \geq 3$): From Vickrey we get $p\mathcal{F}^{(2)}/k$ and from CORE_c we get $(1-p)$ times the quantity in Equation (3) for a combined expected profit of:

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

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. □

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 profit extractor and the latter uses an approximate profit extractor.

We end this section by presenting a pair of results that suggest that there is no straightforward way to improve this result.

6.3 Can we do better?

For basic auctions, CostShare_R is an exact profit extractor for metric $\mathcal{F}$. In this section we show that there is no exact profit extractor for the double auction problem for the natural metrics. The first theorem shows that the profit extractor $\mathcal{F}^{(2)}$-Extract_R that we have presented is best possible, in terms of the fraction α of the profit that is guaranteed to extract. As we have seen, the lack of exact profit extraction is the reason why the competitive ratio we obtain for the double auction problem is worse than the corresponding ratio for the basic auction problem.

Theorem 10. *For input* $\mathbf{b}$ *and* $\mathbf{s}$*, let* $k \geq 2$ *be the number of items sold by* $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$*. Let* $\mathcal{M}_R$ *be a truthful double auction mechanism that is guaranteed to extract a profit of at least* R/α_k *on any* $\mathbf{b}$ *and* $\mathbf{s}$ *with* $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \geq R$*. Then,* $\alpha_k \geq k/(k-1)$*.*

Proof. We prove the lemma assuming the mechanism is deterministic. Since randomized truthful mechanisms are just distributions over deterministic truthful mechanisms, the lemma follows in general from the consideration of this case.

We prove the $k = 2$ case that $\alpha_2 \geq 2$ below. The general theorem follows by a simple, but tedious inductive argument that follows the same lines. For the case $k = 2$, with a bid vector (b_1, b_2, s_1, s_2) consisting of two buyers and two sellers where $b_1 = b_2 = b$ and $s_1 = s_2 = s = b - R/2$. If $\alpha_2 < 2$, then both buyers and both sellers must win. Let $f(\cdot)$ and $g(\cdot)$ denote the bid-independent functions defined by the truthful mechanism $\mathcal{M}_R$. Since all bidders must win, we have that

$$b \geq f(?, b, s, s),\ f(b, ?, s, s) \geq g(b, b, ?, s),\ g(b, b, s, ?) \geq s.$$

Now consider instead the bid vector (b, b', s, s), where $b' > b$. Then since the price offered buyer 2 is the same as it was in the previous case, namely $f(b, ?, s, s)$, and both buyers and sellers must win in order to achieve a revenue which exceeds $R/2$, we still have

$$b \geq f(b, ?, s, s),\ f(?, b', s, s) \geq g(b, b', ?, s),\ g(b, b', ?, s) \geq s.$$

Similarly, if we consider the bid vector (b', b, s, s), we must have

$$b \geq f(?, b, s, s)\ ,\ f(b', ?, s, s) \geq g(b', b, ?, s),\ g(b', b, s, ?) \geq s.$$

Together the facts $b \geq f(b', ?, s, s)$ and $b \geq f(?, b', s, s)$ imply that for the bid vector (b', b', s, s), with $b' > b$, the price offered the buyers is at most b, which in turn implies that the prices paid to the sellers are both at most b.

An analogous argument, starting with a bid vector in which both buyers bid b' and both sellers bid $s' = b' - R/2$ and then moving the seller bids down shows that for any pair of buyers with bid values b' and sellers with bid values $s < b' - R/2$, the sellers both win at prices that are at least $b' - R/2$. This is true in particular for the bid vector (b', b', s, s).

For $b' > b + R/2$, this leads to a contradiction: in the first case, we argued that the prices offered to the sellers were at most b, whereas the second argument shows that the prices offered to the sellers are strictly greater than b. Such prices would give negative profit. □

The final lemma of this section shows that there is no profit extractor, exact or approximate, for the profit metric $\mathcal{F}(\mathbf{b}, \mathbf{s})$.

Lemma 5. *For any value R and $\alpha \geq 1$, there is no truthful mechanism for the double auction problem that achieves a profit of at least R/α on any input $\mathbf{b}$ and $\mathbf{s}$ with $\mathcal{F}(\mathbf{b}, \mathbf{s}) \geq R$.*

Proof. Suppose for a contradiction that such a mechanism $\mathcal{M}_R$ did exist. Consider the single buyer, single seller case with $b_1 = s_1 + R$. On this input, to achieve revenue R/α for any $\alpha \geq 1$, both the buyer and the seller must win the auction. Theorem 1 and $\mathcal{M}_R$'s truthfulness implies that the price for b_1 is given by a bid-independent function $f(s_1)$. Since b_1 wins the auction and the auctioneer has a positive profit, it must be that $f(s_1) \in [s_1, b_1] = [s_1, s_1 + R]$. Symmetrically, we must have $g(b_1) \in [b_1 - R, b_1]$. Now consider inputs $b'_1 = s'_1 + 2R$. Given f and g above, the sell price is in $p'_1 \in [s'_1, s'_1 + R]$ and the buy price is in $q'_1 \in [b'_1 - R, b'_1] = [s'_1 + R, s'_1 + 2R]$. The auctioneer's profit is $p'_1 - q'_1 \leq 0$ which gives a contradiction. □

7 Conclusions

In this paper we have given a game theoretic treatment of the off-line problem of matching up buyers and sellers of a single identical commodity so as to maximize the worst case profit of the arbitrating auctioneer. Open questions related to the double auction problem include considering similar questions in an on-line setting where the buyers, sellers, or both arrive one at a time and the auctioneer must decide whether to sell and at what price before the arrival of the next customer.

The solutions we have presented are based on techniques and ideas that were previously used for the design of competitive basic auctions. We have presented a general reduction that shows how to construct a 2β competitive double auction given a β-competitive basic auction mechanism. We have also presented a generalization of the CORE auction of [11] which can be used in the context of double auctions. This approach gives the best competitive ratio known for the double auction problem. The form in which the CORE mechanism is presented here illustrates how it might be used for mechanism design problems where the goal is profit maximization. Specifically, this approach is constant-competitive against a profit metric $\mathcal{O}$ for auction problems with single-parameter agents for which:

- the profit metric $\mathcal{O}$ is insensitive to minor changes in the relevant inputs;
- there is an approximate profit extractor for the profit metric;
- the profit extractor is safe, in that infeasible outcomes or outcomes with negative profit can not occur.

The fundamental research direction posed by this paper is to gain an understanding of when these conditions hold, i.e. for what kinds of auction problems does the CORE approach work. And specifically, for what problems and metrics do exact or approximate profit extractors exist.

References

1. A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
2. S. Baliga and R. Vohra. Market Research and Market Design. In *Advances in Theoretical Economics*. Volume 3, Issue 1, 2003.
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. A. Blum, V. Kumar, A. Rudra, and F. Wu. Online Learning in Online Auctions. In *Proc. 14th Symposium on Discrete Algorithms*. ACM/SIAM, 2003.
5. A. Blum, T. Sandholm, and M. Zinkevich. Online algorithms for market clearing. In *Proc. 13th Symp. on Discrete Alg.*, pages 971–980, 2002.
6. A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
7. J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
8. E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.
9. 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.
10. Daniel Friedman and John Rust, editors. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison Wesley, 1993.
11. A. V. Goldberg and J. D. Hartline. Competitiveness via Concensus. In *Proc. 14th Symp. on Discrete Alg.* ACM/SIAM, 2003.
12. 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.
13. 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.
14. T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.
15. 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.
16. P. Klemperer. Auction theory: A guide to the literature. In *Journal of Economic Surveys*, pages 227–286. 13(3), 1999.
17. A. Mas-Collel, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
18. R. McAfee. A dominant strategy double auction. In *Journal of Economic Theory*, volume 56, pages 434–450, 1992.
19. H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs. *Economic Theory*, to appear.
20. R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.

21. N. Nisan. Algorithms for Selfish Agents: Mechanism Design for Distributed Computation. In *16th Annual Symp. on Theoretical Aspects of Comp. Sci.*, 1999.
22. 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.
23. N. Nisan and A. Ronen. Computationally feasible VCG Mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.
24. C. Papadimitriou. Algorithms, Games and the Internet. In *Proc. of 33rd Symposium Theory of Computing*, pages 749–752. ACM Press, 2001.
25. T. Sandholm and S. Suri. Market clearability. In *Proc. of the 17th International Joint Conf. on Artificial Intelligence (IJCAI)*, pages 1145–1151, 2001.
26. Combinatorial Auctions: A survey. S. DeVries and R. Vohra. In *Unpublished manuscript*, 2000.
27. W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.
28. 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.