

# Design of a Multi-unit Double Auction E-Market\*

Pu Huang  
Center for Automated Learning and Discovery  
School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213  
phuang@cs.cmu.edu

Alan Scheller-Wolf  
Graduate School of Industrial Administration  
Carnegie Mellon University  
Pittsburgh, PA 15213  
awolf@andrew.cmu.edu

Katia Sycara  
The Robotics Institute  
Carnegie Mellon University  
Pittsburgh, PA 15213  
katia@cs.cmu.edu

---

\*This paper has been accepted for publication in *Computational Intelligence*.

## Abstract

We envision a future economy where e-markets will play an essential role as exchange hubs for commodities and services. Future e-markets should be designed to be robust to manipulation, flexible, and sufficiently efficient in facilitating exchanges. One of the most important aspects of designing an e-market is market mechanism design. A market mechanism defines the organization, information exchange process, trading procedure and clearance rules of a market. If we view an e-market as a multi-agent system, the market mechanism also defines the *structure and rules of the environment* in which agents (buyers and sellers) play the market game. We design an e-market mechanism that is *strategy-proof with respect to reservation price, weakly budget-balanced and individually rational*. Our mechanism also makes sellers unlikely to under-report the supply volume to drive up the market price. In addition, by bounding our market's efficiency loss, we provide fairly unrestrictive sufficient conditions for the efficiency of our mechanism to converge in a strong sense when (1) the number of agents who *successfully trade* is large, or (2) the number of agents, *trading and not*, is large. We implement our design using the RETSINA infrastructure, a multi-agent system development toolkit. This enables us to validate our analytically derived bounds by numerically testing our e-market.

**keywords:** electronic market, auction, mechanism design, multi-agent system, market clearing.

# 1 Introduction

Recent years have seen a growing interest among both academia and industry in Internet marketplaces. Enabled by modern network technologies, e-markets promise nearly friction-free information exchange, broad access to potential buyers and sellers, and almost perfect competition. Launching a new e-market on the Internet is easy nowadays: the underlying software, hardware and network infrastructure are readily accessible. However, this does not mean every e-market can easily survive and succeed: easy launching means that there may exist many similar e-markets competing for potential buyers and sellers. To succeed in this e-market competition, one has to carefully *design* one's e-market to meet the needs of both buyers and sellers, providing attractive features, services, and conveniences for the market players. We envision a future economy where e-markets will play an essential role as exchange hubs for commodities and services. Future e-markets should be robust to manipulation, flexible, and sufficiently efficient in facilitating exchanges. In this paper, we shall show the design a multi-unit double auction (MDA) e-market that is sophisticated enough to satisfy all of these needs.

In an MDA market such as the stock market, sellers and buyers submit "asks" and "bids" respectively. A trade is made if a buyer's bid exceeds a seller's ask. Typically, a seller has multiple units for sale and a buyer wants to purchase more than one unit. Therefore, a seller's ask may match several buyers' bids and a buyer's bid may satisfy several sellers' asks. An MDA e-market must be able to deal with this sort of matching between multiple sellers and multiple buyers involving multiple units. We organize our e-market as an MDA market because it is highly flexible: both buyers and sellers can state how much to trade and at what price. In addition to MDAs, there are other ways to organize an e-market: "take-it-or-leave-it" (prices are set by one side of the market), one-

to-one negotiations, one-side auctions, and single-unit double auctions (SDAs). However, none of these ways of organizing an e-market can accommodate simultaneous trading by multiple buyers and sellers, and at the same time quickly clear a market<sup>1</sup> where each seller sells multiple units and each buyer buys multiple units. Because of their flexibility, we anticipate that MDA e-markets will play an important role in future e-commerce.

The central issue of MDA e-market design is mechanism (trading rule) design. The mechanism of an e-market defines the organization, information exchange process, trading procedure and clearance rules of the market. The market mechanism is also a promise offered by the market maker, who operates and manages the e-market, to all the participating agents. Once the market mechanism is chosen, other related design aspects, such as communication protocol, integration and transaction support, security issues, etc., are implicitly constrained. The communication protocol should be designed to facilitate the information exchange process; transaction support should obey the specific market clearance rules; security measures should be deployed to protect the specific sensitive information, and so on.

The mechanism is announced before the opening of the market so that every agent<sup>2</sup> knows exactly how the market will operate in advance. We anticipate intelligent software agents, acting as delegates of their human masters, becoming important market players in future e-markets. These agents are assumed to be self-interested and autonomous; they pursue their own interests maximizing their own utilities. Taken together they form a multi-agent system. Therefore, we can view an e-market as a multi-agent system where the market mechanism defines the *structure and rules of the environment* in which agents will play the market game. A fundamental issue in mechanism

---

<sup>1</sup>Clearing a market means determining quantities and pairings for all trades.

<sup>2</sup>We use “agents” to refer to both buyers and sellers.

design thus becomes creating an environment where it is in the agents' individual interests to act in a way that also benefits the market as a whole. The MDA mechanism (see Section 3) we designed is *theoretically guaranteed* to ensure this is the case. Specifically, our MDA mechanism is *strategy-proof with respect to reservation price, weakly budget-balanced, asymptotically efficient, and individually rational*. The exact definitions of all these terms will be given in later paragraphs in this section. However, before going into these details, we would like to motivate this research by showing what these properties mean to an individual agent who participates in the e-market. By doing so, we provide an intuitive understanding of the designing goals we are pursuing and why they are important.

Efficiency is one of the most important goals people usually pursue when designing a market mechanism. An efficient market *maximizes* the total profit obtained by all participating agents (Fudenberg and Tirole 1991). An efficient market attracts both buyers and sellers, since it is not biased in favor of either party and it promises to maximize their collective profit. One ideal way of organizing a 100-percent efficient MDA e-market is to let the buyers/sellers submit bids/asks about how many items they want to purchase/sell and at what reservation prices<sup>3</sup>; then based on this information, the market maker solves an optimization problem to determine how many units each agent should purchase/sell and at what price to maximize the total profit<sup>4</sup> of the market (see the problem (LP) in Section 3). However, this method is infeasible because one cannot prevent agents from lying. In fact, the participants in an e-market are usually geographically distributed

---

<sup>3</sup>The reservation price for a buyer is the maximum price he is willing to pay for one unit item; the reservation price for a seller is the minimum price for which she is willing to sell one unit item.

<sup>4</sup>A buyer's profit (or utility) is the quantity he purchases multiplied by the difference between his reservation price and the trading price; a seller's profit is the quantity she sells multiplied by the difference between the trading price and her reservation price.

all over the world; they may be humans or software agents; and the market maker has no control whatsoever over what they report and whether they report truthfully. And most importantly, agents have very strong incentives to lie: every buyer would like to under-report his reservation price and every seller would like to over-report her reservation price. This sort of mis-reporting has two consequences. First, the e-market itself may collapse because the buyers/sellers fear that they will be exploited by the other party. Due to such mistrust, buyers and sellers would exceedingly under-report/over-report their reservation prices, which causes all trades to be impossible and thus hurts all agents. Second, all three parties involved in the e-market, i.e., the market maker, the buyers and the sellers, have no way to determine whether the market is efficient or not, or how efficient it is, because no one has complete information about the market. In fact, if an agent can make an extra profit by cheating, it is more likely that the total loss of all other agents is greater than the extra profit this dishonest agent gets, i.e., the market is more likely to be inefficient. Taken together, this moral hazard, that a dishonest agent could get a higher profit, discourages agents to participate in the e-market and has the potential to ruin the whole e-market.

Obviously, the market maker would like the agents to tell the truth, not only for the reason that it wants to maintain its e-market, but also for the collective interest of all the agents. However, in a free environment like the Internet, we cannot force the agents to reveal the truth. Instead, we need to design a mechanism that induces the agents to tell the truth. This kind of mechanisms is said to be *strategy-proof*, which means that truthfully revealing information is each agent's best strategy independent of what other agents are doing. Strategy-proofness is often achieved by arranging some special form of payments between buyers and sellers. If all these payments sum to zero *exactly* (counting the amount paid by buyers as positive and the amount received by sellers as negative), the mechanism is *exactly budget-balanced*; if they sum to a nonnegative number, the

mechanism is *weakly budget-balanced*. An agent will participate in an e-market only if it expects a nonnegative profit. Therefore, an MDA mechanism must also be designed to be *individually rational*, i.e., it attracts individual agents to voluntarily participate in the e-market because they expect nonnegative *ex ante* profits. Unfortunately, it is well known that *no* mechanism can achieve strategy-proofness, be exactly budget-balanced, efficient and individual rational simultaneously, even for simple one-side auction markets (Myerson and Satterthwaite 1983; Fudenberg and Tirole 1991), let alone our MDA market.

As mentioned before, our MDA mechanism is *strategy-proof with respect to reservation price*, *weakly budget-balanced*, *asymptotically efficient*, and *individually rational*. Strategy-proofness with respect to reservation price means that each agent will honestly report his or her reservation price. In addition to mis-reporting their reservation prices, sellers in an MDA market may also drive up the trading price by mis-reporting the number of units they want to sell (see Section 3). We also consider this problem and show why it is unlikely to happen in our e-market.

Our e-market is asymptotically efficient, which means its efficiency increases to 100% as the number of agents increases. This result is theoretically proved. In fact, it is generally believed that double auctions are highly efficient, especially when the number of agents is large. A number of empirical studies support this belief (Gjerstad and Dickhaut 1998; Cason and Friedman 1996; Plott and Gray 1990; Smith 1982). However, none of these studies showed rigorously how efficiency is achieved in an MDA market. We do so here.

Another property of our e-market is that it is weakly budget-balanced, i.e., there is a surplus after all transactions clear. This surplus is absorbed by the market maker. Clearly, a nonnegative surplus justifies the very existence of the market, otherwise the market needs to be subsidized by outside sources and would not be able to survive for a long time. Weak budget-balance is a relaxation of

exact budget-balance and seems to be a more practical design requirement. In real markets (physical or electronic), market makers are paid for their efforts bringing buyers and sellers together; for example, stock brokers get a certain payment for every transaction they execute for their customers. The market maker of our e-market is paid by the surplus generated by our mechanism.

Many existing e-markets are simple one-side auction markets, such as Ebay and uBid; some laboratory e-markets such as TAC server (Wellman *et al.* 2001) support a double auction mechanism without the guaranteed properties of our mechanism. To the best of our knowledge, our design is the first MDA mechanism that theoretically guarantees all the properties mentioned above.

We implemented our design in a prototype e-market using RETSINA, an open multi-agent infrastructure developed at Carnegie Mellon University (Sycara *et al.* 2002). RETSINA supports quick development of distributed multi-agent systems by providing various system services, reusable common components, agent communication languages, ontology support and management, and visualization tools for developers. Building multi-agent systems using the RETSINA infrastructure, one can concentrate on the specific tasks one's system is supposed to accomplish and enjoy all the services, conveniences and tools provided by RETSINA. RETSINA does not employ any form of centralized control over the multi-agent systems running above it, rather, it presumes that any coordination structure will be implemented by the specific system itself, and provides various facilities to support this. This philosophy ensures that RETSINA is open to various multi-agent system designs. Our e-market design consists of three types of agents: buyers, sellers and a market maker. We incorporated our mechanism into the market maker, developed user-interfaces for the buyer and seller agents and let RETINA handle all the communications and agent interactions.

This paper is organized as follows. Section 2 surveys the related work on double auction research. Papers reviewed are drawn from multiple disciplines, including game theory, experimental eco-



nomics and multi-agent systems. Section 3 shows our mechanism and presents the theorem on strategy-proofness, weakly budget-balance and individual rationality. Section 4 focuses on efficiency. We first define a market efficiency measure and then show how our e-market achieves asymptotic efficiency. Section 5 discusses the implementation issues of our e-market and presents some experimental results. Section 6 concludes this paper.

## **2 Related Work**

Voluminous game theory literature focuses on auction markets, specifically one-side auction markets where there is a monopoly with multiple buyers, or an oligopoly with multiple sellers. However, the literature on double auction market design is limited. Satterthwaite and Williams (1989) and Williams (1991) were among the early researchers studying double auction markets. They designed a single-unit double auction (SDA) market where they simplified the analysis by eliminating the strategic behavior (mis-reporting one's true reservation price) on the sellers' side and showed that the difference between a buyer's bid and his reservation value went to zero in the limit as the number of traders grows. Thus the market converged to efficiency. In Satterthwaite and Williams, no third party was required to balance the market budget. Still in a SDA market, McAfee (1992) allowed strategic behavior on both sides of the market and required a market maker to balance the budget. He proposed a strategy-proof mechanism and showed that the inefficiency converged to zero as the market became large in a weak sense, i.e., the surplus taken by the market maker was not counted as efficiency loss. Barbera and Jackson (1995) characterized a set of strategy-proof mechanism for a multi-unit exchange market where every agent could choose to be a buyer or a seller. Their mechanism was not asymptotically efficient and required a third party to pre-specify

a set of price proposals. Our mechanism extends that of McAfee's in the SDA setting to an MDA market: price is formed collectively by all participating agents instead of being set by a third party, and the efficiency converges in a strong sense as the market grows large.

Another interesting approach to study double auction markets is from a bounded rationality standpoint. Since every agent has limited computational resources, and has to make a decision in a limited time frame, the rationality of every agent is not perfect. Therefore if a mechanism is not strategy-proof it may still be applicable because, due to agents' bounded rationality, they cannot behave strategically (i.e. lie). Gode and Sunder (1993) studied a double auction market full of "zero-intelligence" agents who submitted random bids and asks. They showed that the market was close to efficient even with a few traders. Gjerstad and Dickhaut (1998) allowed agents to use simple rules to form beliefs about their opponents' offers based on trade history and showed that the market price converged to competitive equilibrium quickly. Cason and Friedman (1996) conducted a series of experiments to compare several models by imposing different levels of rationality on agents. Since these studies artificially imposed rationality constraints on trading agents, and it is hard, if not impossible, to model the "upper bound" of rationality of agents in real life auctions, application of these models has been limited.

In recent years, there has been a growing interest in the multi-agent community in auction theory and its application in artificial market design. Babaioff and Nisan (2000) showed how to use double auctions to integrate a decentralized supply chain where each stage in the chain was modeled as an individual agent. Yokoo *et al.* (2001) proposed a double auction mechanism against false-name bids. Das *et al.* (2001) conducted a series of experiments where humans and software agents competed with each other in a double auction market. Sandholm and Suri (2001) showed that if a double auction market allows agents to submit discriminatory bids, the problem of clearing the

market faced by the market maker is NP-Complete, which is reminiscent of the result of Lehmann *et al.* (2000) in one-side combinatorial auction. Several experimental auction systems supporting software agents have also been developed, among them eMediator(Sandholm, 2000) and Auction-Bot (Wurman *et al.* 1998). There is also an annual Trading Agent Competition held as part of main multi-agent conferences, drawing dozens of teams from all over the world to compete (Wellman *et al.* 2001).

### 3 Presentation and Analysis of the Mechanism

In an MDA market with  $m$  buyers and  $n$  sellers, each buyer  $i$  wants to purchase  $X_i$  unit items and each seller  $j$  has  $Y_j$  unit items to sell. We assume both  $X_i$  and  $Y_j$  are public information, i.e., known to every agent. The reservation prices, which are private, for buyer  $i$  and seller  $j$  are  $b_i$  and  $s_j$ , respectively. We assume the reservation price for each agent is static. Let  $q_{ij}$  denote the quantity buyer  $i$  buys from seller  $j$ , and  $p_{ij}$  denote the trading price. Buyer  $i$ 's utility for this transaction is then defined as

$$ub_i = \sum_{j=1}^n (b_i - p_{ij})q_{ij}, \quad (1)$$

and seller  $j$ 's utility is

$$us_j = \sum_{i=1}^m (p_{ij} - s_j)q_{ij}. \quad (2)$$

If all information is public, the maximum total market value, which is the maximum collective utility obtained by all participating agents, can be obtained by solving the linear programming

problem:

$$\begin{aligned}
& \max \sum_{i=1}^m \sum_{j=1}^n q_{ij} (b_i - s_j) \\
& \sum_{i=1}^m q_{ij} \leq Y_j \quad \text{for every seller } j \\
& \sum_{j=1}^n q_{ij} \leq X_i \quad \text{for every buyer } i \\
& q_{ij} \geq 0 \quad \forall i, j
\end{aligned} \tag{LP}$$

The two constraints state that a seller sells no more than what she possesses and a buyer will not buy more than he needs. The trading price does not show up in the problem (LP) because the maximum total market value is determined collectively by each agent's reservation price. If buyer  $i$  buys quantity  $q_{ij}$  from seller  $j$  at price  $p_{ij}$ , then the market value this transaction implements is the sum of buyer  $i$ 's utility plus seller  $j$ 's utility, which is  $(b_i - s_j)q_{ij}$ , independent of  $p_{ij}$ , the trading price. However, it is clear that the trading price will affect *each* agent's utility.

In most situations, the centralized solution is not available because agents have their own private reservation prices. To induce the agents to report their true reservation prices, we apply a Vickrey-like auction on each side of the market. On the buyers' side, each buyer  $i$  reports a price  $rb_i$  (which may or may not equal  $b_i$ ); and on the sellers' side, each seller  $j$  reports a price  $rs_j$  (which may or may not equal  $s_j$ ). Without loss of generality, we assume

$$rb_1 > rb_2 \dots > rb_m, \tag{3}$$

and

$$rs_1 < rs_2 \dots < rs_n. \tag{4}$$

We assume strict order relations here because if two buyers report the same reservation price, we can add their volumes together to form an equivalent bid. The same thing may be done with the sellers. Agents are allowed to split their volumes. For example, a buyer who wants to buy 6 items

may submit a bid of \$20 for 2 items and \$60 for the remaining 4 items. In this case, this buyer will be treated as two different buyers. However, as will become clear in the following paragraphs, as we assume every agent has only one reservation price, he/she will always submit a bid or a ask at his/her reservation price and will not split the volume.

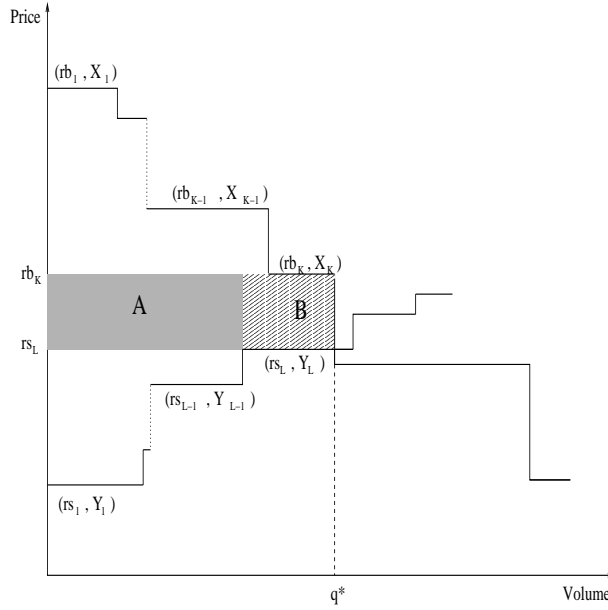


Figure 1: A multi-unit double auction market where  $K - 1$  buyers and  $L - 1$  sellers successfully trade.

Our mechanism works as follows. It arranges the demand volumes according to the ascendent price order as shown in (3) and the supply volumes according to the decedent price order as shown in (4) (refer to Figure 1). At the critical point  $q^*$  where the aggregate demand and supply meet, there are a buyer  $K$  and a seller  $L$ . Either their reported prices satisfy

$$rb_K \geq rs_L \geq rb_{K+1}, \tag{5}$$

and the aggregate demand and supply satisfy

$$\sum_{j=1}^{L-1} Y_j \leq \sum_{i=1}^K X_i \leq \sum_{j=1}^L Y_j \tag{6}$$

(Case I, as shown in Figure 1), or their reported prices satisfy

$$rs_{L+1} \geq rb_K \geq rs_L \quad (7)$$

and the aggregate demand and supply satisfy

$$\sum_1^{K-1} X_i \leq \sum_1^L Y_j \leq \sum_1^K X_i, \quad (8)$$

(Case II).

We state our mechanism for Case I; Case II is similar. To clear the market, we first check whether inequality

$$\sum_1^{K-1} X_i \geq \sum_1^{L-1} Y_j, \quad (9)$$

or

$$\sum_1^{K-1} X_i \leq \sum_1^{L-1} Y_j, \quad (10)$$

holds. If (9) holds, as in Figure 1, we follow

**rule1:** all the sellers with indices  $j < L$  sell all their volume  $Y_j$  at price  $rs_L$ ; all the buyers with indices  $i < K$  trade at price  $rb_K$  and each of them buys a volume equal to  $X_i - (\sum_1^{K-1} X_i - \sum_1^{L-1} Y_j)/(k-1)$ <sup>5</sup>.

If (10) holds, we follow

**rule2:** all the buyers with indices  $i < K$  buy all their volume  $X_i$  at price  $rb_K$ ; all the sellers with indices  $j < L$  trade at price  $rs_L$  and each of them sells a volume equal to  $Y_j - (\sum_1^{L-1} Y_j - \sum_1^{K-1} X_i)/(L-1)$ .

---

<sup>5</sup>If  $(\sum_1^{K-1} X_i - \sum_1^{L-1} Y_j)/(k-1) > X_i$  for some buyer  $i$ , this buyer buys nothing and the “burden left”,  $(\sum_1^{K-1} X_i - \sum_1^{L-1} Y_j)/(k-1) - X_i$ , is averaged over the  $K-2$  buyers left. Continue this procedure until each buyer left trades a positive volume. A similar procedure may also apply to **rule2**.

If (9) holds, there is a over-demand, and this over-demand is averaged over all the first  $K - 1$  buyers; If (10) holds, there is a over-supply, and this over-supply is split over all the first  $L - 1$  sellers.

The total trade volume is

$$\min\left(\sum_1^{K-1} X_i, \sum_1^{L-1} Y_j\right).$$

The trading surplus  $(rb_K - rs_L) \min(\sum_1^{K-1} X_i, \sum_1^{L-1} Y_j)$  (shaded region  $A$  in Figure 1) is taken by the market maker. The potential trading value represented by the shaded region  $B$  is sacrificed, not collected by any of the three parties, the sellers, the buyers or the market maker, involved in the market. Since either buyers or sellers are forced to sacrifice part of their volumes, there is yet another potential trading value  $C$  lost (not shown in the figure). If (9) holds, this loss is bounded by

$$C \leq (rb_1 - rb_K)Y_L, \quad (11)$$

and if (10) holds, this loss is bounded by

$$C \leq (rs_L - rs_1)X_K. \quad (12)$$

Considering (9) and (6) (or (9) and (8)) together, we get  $\sum_1^{K-1} X_i - \sum_1^{L-1} Y_j \leq Y_L$ . Since  $rb_1 - rb_K$  is the maximum price margin for buyers, inequality (11) holds; Considering (10) and (6) (or (10) and (8)) together, we get  $\sum_1^{L-1} Y_j - \sum_1^{K-1} X_i \leq X_K$ . Since  $rs_L - rs_1$  is the maximum price margin for sellers, inequality (12) holds. Note that part  $C$ , as part  $B$ , is also sacrificed and not collected by any of the three parties involved in the market.

To summarize, our mechanism is as follows:

**Mechanism:** The market maker first sorts the reported reservation prices from the buyers and sellers according to inequalities (3) and (4). Every buyer with index less than  $K$  and every seller with index less than  $L$  will trade. Then depending on whether the market situation is Case I or Case II, and whether inequality (9) or (10) holds, the market maker decides whether to apply **rule1** or **rule2**.

The above mechanism is announced in advance, and every agent knows that these rules will be followed by the market maker to operate the e-market. What this means to an individual agent is stated in the following theorem.

**Theorem 1:** Under the assumption that the buyers and sellers' volumes are public information, the above mechanism is strategy-proof with respect to reservation price, weakly budget-balanced, and individually rational.

**Proof:** We'll prove for the scenario where inequality (12) holds. In this scenario, we apply **rule2**. The proof of strategy-proofness for the buyers is the same as Vickrey's argument. Suppose a buyer  $i$  with reservation price  $b_i$  reports  $rb_i$ . If  $b_i \geq rb_K$ , over-bidding, i.e.,  $rb_i > b_i$ , will give the buyer the same utility as if he bids  $b_i$ ; under-bidding, i.e.,  $rb_i < b_i$ , may cause him to lose the trade he otherwise would win, and even if he wins the trade, he gets the same utility as if he bids  $b_i$ . If  $b_i < rb_K$ , over-bidding may incur negative utility if the buyer is included in the final trade. Even if the buyer is not included, he still gets zero, the same as if he bids  $b_i$ ; under-bidding just gives him utility zero, the same as if he bids  $b_i$ .

In scenario (12), sellers who successfully trade only sell part of their volumes. However, since every seller who trades takes the same volume decrease, and no seller with index  $j < L$  can mitigate her trading volume decrease by mis-reporting her true reservation price, for the same reasons as shown for the buyers, every seller will report her true reservation price.



The proof for the scenario (11) is similar.

It is clear that  $b_K \geq s_L$ , so the mechanism always gets nonnegative payment and hence it is weakly budget-balanced.

Since every agent gets positive utility if it trades, or zero if it does not, our mechanism is individually rational.  $\square$

A minor change of our mechanism will make it exactly budget-balanced. If after balancing the demand and supply by applying **rule1** or **rule2**, we set the trading price to be any number between  $rb_K$  and  $rs_L$  ( buyer  $K$  and seller  $L$  do not trade), then the total payments from all agents add up to zero and the market maker gets nothing. However, as stated before, we consider weak budget-balance a more proper design goal since the market maker should be paid for managing the e-market.

As showed in the proof of Theorem 1, the primary reason for our mechanism being strategy-proof is that everyone who trades always pays (or sells at) the price proposed by someone else. Actually, this is also a general principle to follow when designing a strategy-proof mechanism. In our MDA market, we also need to match the total demand volume and the total supply volume exactly. To do so, and at the same time keep our mechanism strategy-proof, we sacrifice the potential trading value  $B$  and  $C$ .

Strategy-proofness of a mechanism usually requires some specific utility function form. For example, Vickrey's mechanism is strategy-proof when each agent's utility function is *quasi-linear* (Vickrey 1961; Clarke 1971; Groves 1973). Similarly, our mechanism requires that the trading volume and price ( $q_{ij}$  and  $p_{ij}$  in definition (1) and (2)) of each agent are "separable" in its utility function, which simply means that  $p_{ij}$  does not depend on  $q_{ij}$ . If  $p_{ij}$  and  $q_{ij}$  are related, for example, a seller may offer a discount price if a buyer wants to buy a large volume (discriminatory bidding),

then complicated strategic mis-reporting of reservation price may be possible. Another problem related to discriminatory bidding is that clearing the market becomes very difficult if there is a large number of agents. Actually, clearing the market in discriminatory bidding is NP-Complete (Sandholm and Suri 2001). Since we do not allow discriminatory bidding, our market clearance algorithm is simple. It can be done either by sorting the reported prices and searching for the agents “at the edge”, as done by our mechanism, or by solving a linear programming problem as shown in (LP). However, note that the (LP) problem is under-constrained and has multiple optimal solutions, one needs to find the unique solution corresponding to our mechanism.

We assume the volume of each agent is public information; in other words, each agent cannot misrepresent his/her volume. In a real-life market, buyers usually will honestly report their required volumes and seek low prices; a seller however has the incentive to *under-report* the quantities she possesses in order to tighten the supply and drive the market price up. This can be reflected in our market as follows (refer to Figure 1). A seller  $j < L$  may under-report her volume to shift the supply curve to the left. By doing so, she hopes to put the seller  $L + 1$  “on the edge”. If every seller who trades pays the price reported by the seller “at the edge”, then seller  $j$  may be better off using this strategy: she may get a higher utility by selling a lower volume at the higher price instead of a larger volume at the lower price. Because of the presence of such sophisticated strategies, exact analysis of MDA markets where volumes are also private information is extremely difficult. For an example, refer to Figure 2, which is a copy of Figure 1. In Figure 2, we assume all the other bids and asks are fixed and only seller 1 under-reports her volume. If seller 1 honestly reports her true volume, she will get a utility equal to the area of the rectangle  $FH$ . If she under-reports by an amount  $d$ , she puts seller  $L + 1$  on the edge. Seller 1 then gets a higher utility equal to the area of the rectangle  $FI$ . If she continues under-reporting the volume by another amount  $e$ , seller  $L + 2$  moves

to edge, and seller 1 gets a utility equal to the area of rectangle  $FJ$ , which is smaller than the area of the rectangle  $FH$ . In Figure 3, we show the utility of seller 1 as a function of her own reported volume. Point  $T$  represents seller 1 reporting her true volume; Point  $S$  represents the utility she gets by under-reporting by an amount  $d$ ; Point  $W$  represents the utility of under-reporting by an amount  $d + e$ . In this example, seller 1 gets a lower utility in point  $W$  than in the original point  $T$ . It should be noted that the relation  $W < T < S$  is for this example only, and will not hold in general.

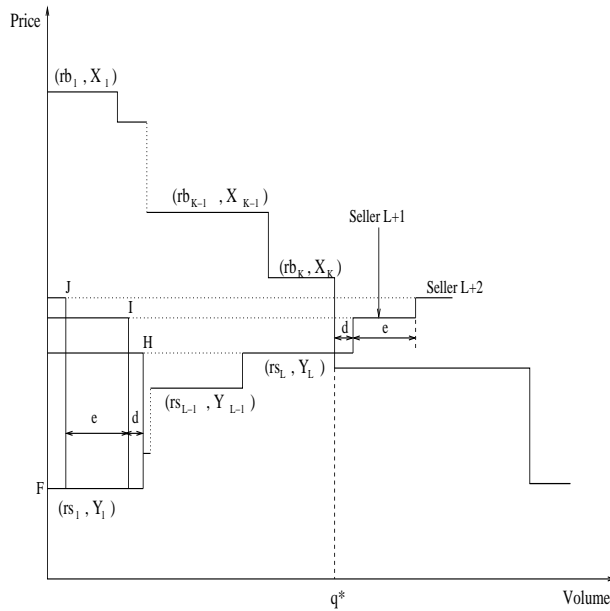


Figure 2: Seller 1 under-reports her volume, assuming all other bids and asks in the market are fixed.

Although possible, this strategy of under-reporting is difficult for the sellers to successfully implement in our e-market for several reasons. First, only sellers included in the final trade ( $j < L$ ) can use this strategy to affect the trading price, but an individual seller does not know whether she will be included or not when submitting an ask. Second, a seller cannot decide how much to under-report without full information of the whole market; and arbitrarily under-reporting may




Figure 3: Seller 1's utility as a function of her reported volume, assuming all other bids and asks in the market are fixed.

actually lower her utility, as shown in Figure 3. Third, even if all sellers had full information about the whole market (though our market mechanism does not provide them this information), they would still have a difficult time deciding how much to under-report, because all sellers have the incentive to under-report and an individual seller's decision must take the others' decisions into consideration. Thus, every seller faces an iterative induction game: she knows everyone may mis-report, everyone knows she knows everyone may mis-report, ....., and so on and so forth *ad infinitum*. This problem structure is called "common knowledge" in game theory. It is generally exceedingly difficult to find equilibrium solutions to this kind of game, especially when there are multiple players.

In general, if all buyers report honestly, then any seller's utility is a multi-variable function over all the sellers' reported volumes. We can view this function as a rugged surface in a high dimensional space. Figure 3 is the projection of this surface into one dimension. In this surface, point  $T$  is a local maximum. We posit that the lack of further information makes any seller unlikely to deviate from  $T$  and explore other parts of the surface.

## 4 Efficiency

Traditionally, economists view efficiency as a measure of how much of the maximum market value a mechanism can “induce”, whereas how the implemented value is distributed is not a concern. For instance, in a classical single-unit Vickrey auction, the maximum total market value equals the highest bidder’s reservation price  $v_1$ . The highest bidder gets the item at the second highest bidders’ reservation price  $v_2$  and his utility is hence  $v_1 - v_2$ . All the other bidders get 0 and the seller gets  $v_2$ . The collective utility obtained by all agents is thus  $v_1$ , so we say that the Vickrey mechanism is efficient because it implements the maximum total market value. In a double auction market, when counting the total market value a mechanism can implement, the surplus taken by the market maker is usually included, as in McAfee (1992), even though this surplus is not taken by any participating agents. Under this definition, McAfee showed that a *single-unit* double auction is asymptotically efficient. We show that our mechanism is asymptotically efficient for *multi-unit* double auctions under this definition. We achieve this by showing that even under a stronger efficiency definition, our mechanism is still asymptotically efficient. By stronger efficiency definition, we mean that the efficiency is defined as a measure of how much all traders actually obtain compared to the maximum total market value: The surplus taken by the market maker is considered as an efficiency loss. We refer to the traditional efficiency definition as weak and ours as strong. From the participating agents’ point of view, our efficiency definition makes more sense: participating agents are only concerned with how much *they* obtain; the surplus taken by the market maker actually is a loss to them. We show that even under this stronger definition, asymptotic efficiency still holds for our mechanism.

Assuming  $X_i$ ,  $i = 1, \dots, m$ , the demand volumes from the  $m$  buyers, are independent random

samples equal in distribution to a random variable  $X$ ; and  $Y_j$ ,  $j = 1, \dots, n$ , the supply volumes from the  $n$  sellers, are independent random samples equal in distribution to a random variable  $Y$ . We further assume the reservation prices of the  $m$  buyers are drawn independently from a distribution  $F$  with continuous density  $f$  and support on the compact interval  $[\underline{b}, \bar{b}]$ ; and the reservation price of the  $n$  sellers are drawn independently from a distribution  $G$  with continuous density  $g$  and support on the compact interval  $[\underline{s}, \bar{s}]$ . We also require that the density functions  $f$  and  $g$  have nonzero minimum values, i.e.,

$$\phi = \min\{f(x) : \underline{b} \leq x \leq \bar{b}\} > 0, \quad (13)$$

$$\gamma = \min\{g(x) : \underline{s} \leq x \leq \bar{s}\} > 0. \quad (14)$$

An illustration of distribution functions satisfying these assumptions can be found in Figure 4. These assumptions are sufficient to show efficiency under the weak definition in a simple SDA market (McAfee 1992). We shall show that these assumptions are sufficient to show the efficiency under the strong definition in an MDA market if the number of agents who successfully trade is large (Theorem 2).

If  $K$  among the  $m$  buyers and  $L$  among the  $n$  sellers successfully trade, then, as stated in Section 3, the total efficiency loss includes three parts: the surplus taken by the market maker (the shaded region A in Figure 1), the sacrificed trading value between buyer  $K$  and seller  $L$  (the shaded region B in Figure 1), and the sacrificed trading value  $C$  (as shown in inequality (11) or (12)). Define the total efficiency loss as

$$t(K, L) = A + B + C = (b_K - s_L)q^* + C.$$

The total market value can be written as

$$v(K, L) = \sum_1^K (b_i - b_K) X_i + \sum_1^L (s_L - s_j) Y_j + (b_K - s_L) q^*.$$

Define

$$ir(K, L) = \frac{E[t(K, L)]}{E[v(K, L)]}$$

as the market “inefficiency ratio ”given  $K$  and  $L$  are known. Since both agents’ reservation prices and volumes are random variables, we take expectations of  $t(K, L)$  and  $v(K, L)$  and use their ratio as a measure of how much efficiency is lost compared to the total market value. Our first step is to show that  $ir(K, L)$  converges to zero as  $K$  and  $L$  become large. We need the following lemma, which says the difference between the reservation prices of consecutive buyers/sellers always has the same order of magnitude as the inverse of the number of buyers/sellers.

**Lemma 1:**

$$\frac{1}{\varphi(m+1)} \leq E[b_i - b_{i+1}] \leq \frac{1}{\phi(m+1)}, \quad (15)$$

for all  $i = 1, 2, \dots, m-1$ ; and

$$\frac{1}{\delta(n+1)} \leq E[s_{j+1} - s_j] \leq \frac{1}{\gamma(n+1)}. \quad (16)$$

for all  $j = 1, 2, \dots, n-1$ , where  $\phi$  and  $\gamma$  are constants defined in (13) and (14), and  $\varphi$  and  $\delta$  are constants defined as

$$\varphi = \max\{f(x) : \underline{b} \leq x \leq \bar{b}\} \geq \phi > 0, \quad (17)$$

$$\delta = \max\{g(x) : \underline{s} \leq x \leq \bar{s}\} \geq \gamma > 0. \quad (18)$$

**Proof:** We prove the case for  $E[s_{j+1} - s_j]$ . The case for the buyers is similar.

The joint distribution of  $s_j$  and  $s_{j+1}$  is

$$g_{j,j+1}(x, y) = n(n-j)C_{n-1}^{j-1}G(x)^{j-1}g(x)g(y)(1-G(y))^{n-j-1},$$

where  $C_{n-1}^{j-1} = \frac{(n-1)!}{(j-1)!(n-j)!}$  (refer to, e.g., David 1970). We can calculate the expectation directly:

$$\begin{aligned} & \frac{1}{n(n-j)C_{n-1}^{j-1}} E[s_{j+1} - s_j] \\ &= \int_{\underline{s}}^{\bar{s}} \int_{\underline{s}}^y (y-x)G(x)^{j-1}g(x)g(y)(1-G(y))^{n-j-1} dx dy \\ &= \int_{\underline{s}}^{\bar{s}} g(y)(1-G(y))^{n-j-1} \times \left[ \frac{(y-x)G(x)^j}{j} \Big|_{\underline{s}}^y + \int_{\underline{s}}^y \frac{G(x)^j}{j} dx \right] dy \\ &= \int_{\underline{s}}^{\bar{s}} \frac{g(y)(1-G(y))^{n-j-1}}{j} \left[ \int_{\underline{s}}^y G(x)^j dx \right] dy. \end{aligned}$$

Define

$$U(y) = \int_{\underline{s}}^y G(x)^j dx,$$

then the above integral can be written as

$$\begin{aligned} & \int_{\underline{s}}^{\bar{s}} \frac{g(y)(1-G(y))^{n-j-1}U(y)}{j} dy \\ &= \frac{-U(y)(1-G(y))^{n-j}}{j(n-j)} \Big|_{\underline{s}}^{\bar{s}} + \int_{\underline{s}}^{\bar{s}} \frac{(1-G(y))^{n-j}}{j(n-j)} dU(y) \\ &= \int_{\underline{s}}^{\bar{s}} \frac{G(y)^j(1-G(y))^{n-j}}{j(n-j)} dy. \end{aligned}$$

so, we have

$$E[s_{j+1} - s_j] = C_n^j \int_{\underline{s}}^{\bar{s}} G(y)^j (1-G(y))^{n-j} dy. \quad (19)$$

Note that by the integral transformation  $u = G(y)$ , we can write  $E[s_{j+1} - s_j]$  as

$$E[s_{j+1} - s_j] = C_n^j \int_0^1 g^{-1}(u) u^j (1-u)^{n-j} du,$$



where  $g^{-1}(u) = \frac{dG^{-1}(u)}{du}$  is the derivative of the inverse distribution function  $G^{-1}(u)$  and

$$g^{-1}(u) = \frac{1}{g(y)}.$$

Since  $g(y)$  is continuous on a closed interval  $[\underline{s}, \bar{s}]$ , it has a maximum value  $\delta = \max\{g(y) : \underline{s} \leq x \leq \bar{s}\}$ . Using assumption (14), we have

$$\frac{1}{\delta} \leq g^{-1}(u) \leq \frac{1}{\gamma}.$$

Noting that

$$\int_0^1 u^j (1-u)^{n-j} du = \frac{j!(n-j)!}{(n+1)!},$$

we have

$$\frac{1}{\delta(n+1)} \leq E[s_{j+1} - s_j] \leq \frac{1}{\gamma(n+1)}. \quad \square$$

Equation (19) was first found by Pearson (1902) a century ago. Lemma 1 can be directly applied to show the main result in McAfee (1992) in a much simpler manner than presented there.

Given Lemma 1, the following theorem is straight-forward, which says that the efficiency converges linearly at a speed proportional to the number agents who successfully trade.

**Theorem 2:** Assume  $0 < E[X], E[Y] < \infty$ , then

$$ir(K, L) \leq \Delta \max\left(\frac{1}{K-1}, \frac{1}{L-1}\right), \quad (20)$$

where  $\Delta$  is a constant.

**Proof:** Rewrite  $v(K, L)$  as

$$\begin{aligned} v(K, L) = & (b_1 - b_2)X_1 + (b_2 - b_3)(X_1 + X_2) + \dots + (b_{K-1} - b_K) \sum_1^{K-1} X_i \\ & + (s_2 - s_1)Y_1 + (s_3 - s_2)(Y_1 + Y_2) + \dots + (s_L - s_{L-1}) \sum_1^{L-1} Y_j \\ & + (b_K - s_L)q^*. \end{aligned}$$

By Lemma 1,

$$\begin{aligned}
E[v(K, L)] &\geq \frac{1}{\varphi(m+1)} E[X_1 + (X_1 + X_2) + \dots + (\sum_1^{K-1} X_i)] \\
&\quad + \frac{1}{\delta(n+1)} E[Y_1 + (Y_1 + Y_2) + \dots + (\sum_1^{L-1} Y_j)] \\
&= \underbrace{\frac{K(K-1)E[X]}{2\varphi(m+1)}}_c + \underbrace{\frac{L(L-1)E[Y]}{2\delta(n+1)}}_d.
\end{aligned}$$

We prove the theorem for CASE I. In CASE I,  $q^* = \sum_1^K X_i$ , inequality (5) holds and hence

$$E[t(K, L)] \leq E[b_K - b_{K+1}] E[\sum_1^K X_i] + E[C].$$

If (11) holds,  $E[t(K, L)]$  can be further bounded as

$$\begin{aligned}
E[t(K, L)] &\leq E[b_K - b_{K+1}] E[\sum_1^K X_i] + E[(b_1 - b_K)Y_L] \\
&\leq \frac{KE[X]}{\phi(m+1)} + \frac{KE[Y]}{\phi(m+1)},
\end{aligned}$$

and therefore

$$\begin{aligned}
ir(K, L) &\leq (\frac{KE[X]}{\phi(m+1)} + \frac{KE[Y]}{\phi(m+1)}) / (c + d) \\
&\leq (\frac{KE[X]}{\phi(m+1)} + \frac{KE[Y]}{\phi(m+1)}) / c \tag{21} \\
&= \frac{1}{K-1} (\frac{2\varphi}{\phi} + \frac{2\varphi E[Y]}{\phi E[X]}).
\end{aligned}$$

Since  $E[X]$  and  $E[Y]$  are constants,  $ir(K, L)$  is bounded by  $1/(K-1)$  multiplied by a constant.

If (12) holds, we can bound  $E[t(K, L)]$  as

$$\begin{aligned}
E[t(K, L)] &\leq E[b_K - b_{K+1}] E[\sum_1^K X_i] + E[(s_L - s_1)X_K] \\
&\leq \underbrace{\frac{KE[X]}{\phi(m+1)}}_a + \underbrace{\frac{LE[X]}{\gamma(n+1)}}_b.
\end{aligned}$$

Note that  $\frac{a+b}{c+d} \leq \max(\frac{a}{c}, \frac{b}{d})$ , therefore in this scenario we have

$$ir(K, L) \leq \max(\frac{1}{K-1} \frac{2\varphi}{\phi}, \frac{1}{L-1} \frac{2\delta E[X]}{\gamma E[Y]}). \tag{22}$$

Combining inequalities (21) and (22), we can bound  $ir(K, L)$  as in inequality (20). The proof for CASE II is similar.  $\square$

Theorem 2 says that if the number of agents who *successfully trade* is large, the market inefficiency converges to zero at rate  $\max(1/(K-1), 1/(L-1))$ . However, this theorem does not answer the question of whether the market efficiency increases as the number of agents, *trading and not*, increases. To measure the market efficiency in terms of  $m$  and  $n$  instead of  $K$  and  $L$ , we define the expected market inefficiency ratio as:

$$IR(m, n) = \sum_{K=1}^m \sum_{L=1}^n P(K, L) ir(K, L), \quad (23)$$

where  $P(K, L)$  denotes the probability that there are  $K$  buyers and  $L$  sellers successfully trading given that there are  $m$  and  $n$  buyers and sellers in the market respectively. As the number of agents becomes large, we show that

$$IR(\infty, \infty) = \lim_{m, n \rightarrow \infty} IR(m, n) \rightarrow 0.$$

**Theorem 3:** A sufficient condition for market inefficiency going to zero is that  $K$  and  $L$  go to infinity as  $m$  and  $n$  go to infinity.

**Proof:** Let  $P(K = \infty, L = \infty)$  denote the probability that  $K$  and  $L$  go to infinity as  $m$  and  $n$  go to infinity. We thus assume this probability goes to 1, i.e.,

$$\lim_{m, n \rightarrow \infty} P(K = \infty, L = \infty) \rightarrow 1. \quad (24)$$

Let  $P(K \leq m_0, L \leq n_0)$  denote the the probability that  $K \leq m_0$  and  $L \leq n_0$ . We have

$$\begin{aligned} & IR(\infty, \infty) \\ &= \sum_{K=1}^{m_0} \sum_{L=1}^{n_0} P(K, L) ir(K, L) + \sum_{K=m_0+1}^{\infty} \sum_{L=1}^{n_0} P(K, L) ir(K, L) \end{aligned}$$

$$\begin{aligned}
& + \sum_{K=1}^{m_0} \sum_{L=n_0+1}^{\infty} P(K, L)ir(K, L) + \sum_{K=m_0+1}^{\infty} \sum_{L=n_0+1}^{\infty} P(K, L)ir(K, L) \\
\leq & P(K \leq m_0, L \leq n_0) + P(K \geq m_0 + 1, L \leq n_0) \\
& + P(K \leq m_0, L \geq n_0 + 1) + \Delta \max\left(\frac{1}{m_0}, \frac{1}{n_0}\right)P(K \geq m_0 + 1, L \geq n_0 + 1)
\end{aligned}$$

The first three terms of the r.h.s. of the above inequality are obtained by observing that

$$ir(K, L) \leq 1, \forall K, L.$$

The fourth term is obtained by using the fact that

$$ir(K, L) \leq \Delta \max\left(\frac{1}{m_0}, \frac{1}{n_0}\right), \forall K \geq m_0 + 1, L \geq n_0 + 1,$$

which in turn can be derived from inequality (20) in Theorem 2.

As  $m$  and  $n$  go to infinity, given any  $\epsilon > 0$ , we can always find finite  $m_0$  and  $n_0$  such that

$$\Delta \max\left(\frac{1}{m_0}, \frac{1}{n_0}\right) \leq \epsilon/2.$$

Hence the fourth term in the r.h.s. of the above inequality is bounded by  $\epsilon/2$ , as  $P(K \geq m_0 + 1, L \geq n_0 + 1) \leq 1$ . As  $m$  and  $n$  become large, the sum of the first three terms can be bounded by  $\epsilon/2$  as well, implied by assumption (24). Therefore,  $IR(\infty, \infty) \rightarrow 0$  as  $m$  and  $n$  go to infinity.  $\square$

Condition (24) is very easy to satisfy. If we restrict each agent's trade volume to have finite expectation, then as long as  $[\underline{b}, \bar{b}]$  and  $[\underline{s}, \bar{s}]$  intersect, condition (24) holds. Figure 4 shows one possible way  $[\underline{b}, \bar{b}]$  and  $[\underline{s}, \bar{s}]$  can intersect, which is  $\underline{b} \leq \underline{s} \leq \bar{b} \leq \bar{s}$ . We show how this implies (24). Choose  $u$  such that  $[\underline{s} < u < \bar{b}]$  then  $u$  divides the interval  $[\underline{s}, \bar{b}]$  into two parts,  $[\underline{s}, u]$  and  $[u, \bar{b}]$ . The number of sellers whose reservation prices fall into  $[\underline{s}, u]$  goes to infinity as  $n$  goes to infinity; and the number of buyers whose reservation prices fall into  $[u, \bar{b}]$  goes to infinity as  $m$  goes to infinity. Further, either the number of buyers who trade is greater than the number of buyers in  $[u, \bar{b}]$ , or the




Figure 4: One possible way  $[\underline{b}, \bar{b}]$  and  $[\underline{s}, \bar{s}]$  intersect.

number of sellers who trade is greater than the number of sellers in  $[\underline{s}, u]$ . Assume the former. In this case, we already know the number of buyers who trade goes to infinity, and need only show that the number of seller who trade goes to infinity. If the number of buyers who trade goes to infinity, then the total trade volume must go to infinity as well. This total volume is sold by the sellers who successfully trade. Since every seller's volume has finite expectation, there must be infinite number of sellers to support this total volume. Therefore the number of sellers who trade goes to infinity as well. Thus condition (24) is satisfied.

There are three other ways  $[\underline{b}, \bar{b}]$  and  $[\underline{s}, \bar{s}]$  may intersect. They are  $\underline{s} \leq \underline{b} \leq \bar{s} \leq \bar{b}$ ,  $\underline{s} \leq \underline{b} \leq \bar{b} \leq \bar{s}$ , and  $\underline{b} \leq \underline{s} \leq \bar{s} \leq \bar{b}$ . The proof that for these cases condition (24) holds is similar. Note that if  $\bar{b} < \underline{s}$ , no trade is possible.

## 5 Implementation and Experiments

RETSINA provides a set of Agent Foundation Classes supporting various agent types. We use three types of them: task agents, interface agents, and middle agents. The buyers and sellers are imple-

mented as task agents; the market maker is derived from the default middle agent of RETSINA, the MatchMaker; and we also use the standard interface agent, DemoDisplay in RETSINA to visualize and monitor our e-market (RETSINA AFC developers Guide 2002).

Task agents in RETSINA specialize in some specified tasks they are supposed to accomplish. In our e-market, seller and buyer agents only need to submit asks and bids respectively, and accept the information sent by the market maker about how much they sell or purchase and at what price. This information exchange is handled by the communication component provided by RETSINA (Shehory and Sycara, 2000). We show a seller agent in Figure 5; it has two input boxes to submit asks and two output boxes to accept information from the market maker. Buyer agents are similar.

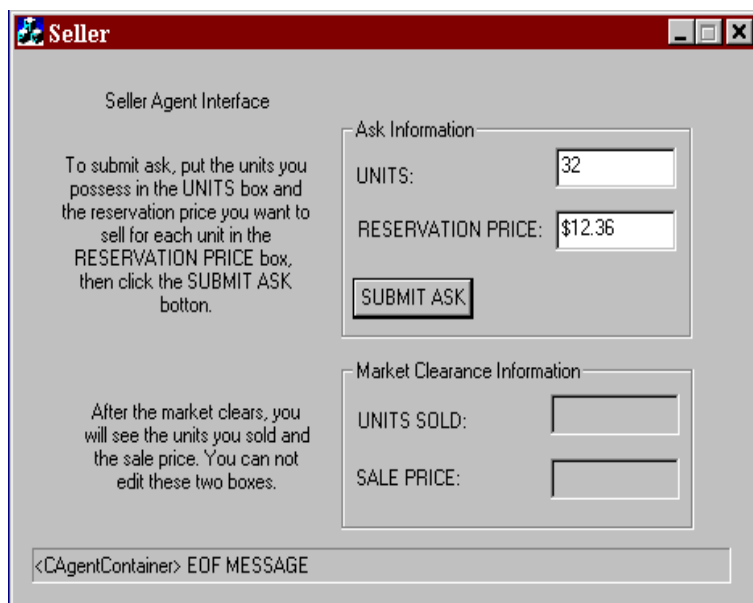


Figure 5: A seller agent interface.

The role of the market maker is to bring demand and supply together and execute the designed mechanism to clear the e-market. This role fits perfectly with the role played by the RETSINA MatchMaker, which serves as a “yellow pages” of agent capabilities, matching service providers with service requesters based on agent capability descriptions. Our market maker is derived from

the MatchMaker, and hence inherits the communication capacity from it. Our designed mechanism instructs the market maker to match the demand from the buyer agents with the supply provided by the seller agents. After the e-market clears, the market maker needs to send messages to all participating agents, as mentioned before. To find the specific location of each agent to send a message, the market maker invokes the Agent Name to Location Mapping Service (ANS) in RETSINA. The location of a participating agent may be anywhere over the Internet; this stand-alone ANS increases the system flexibility.




Figure 6: A MDA e-market with three buyers and three sellers. Lines linking the agents and the MarketMaker represent that they are communicating. These lines are dynamically displayed as the agents submit asks and bids.

Figure 6 visualizes our e-market using the RETSINA DemoDisplay. Each agent in the DemoDisplay is represented by a small icon; lines linking them represent that they are exchanging messages. Our e-market also keeps recording logs of every activity that takes place and every transaction executed.

We let our agents generate random bids and asks to test our e-market. We study two types of reservation price distributions, one is a uniform, the other is a beta. A uniform price distribution represents the case where agents' prices are equally distributed over the support, while a beta distribution represents the case where agents' prices aggregate in some parts of the support. We assume both the buyers' volume  $X_i$  and the sellers' volume  $Y_j$  are drawn from a uniform distribution  $U(0, 1)$ .




Figure 7:  $IR(m, n)$  as a function of  $m$  and  $n$  when agents' reservation prices are uniformly distributed.

In the first experiment, we set the sellers' reservation price distribution to be uniform  $g = U(1, 2)$ , and shift the buyers' reservation price distribution from  $f_1 = U(0.5, 1.5)$  to  $f_2 = U(1, 2)$  to  $f_3 = U(1.5, 2.5)$ . For each distribution pair  $(f_i, g)$ , we run 8 simulations, each of which corresponds to a certain number of participating agents. (We assume the number of buyers equals the number of sellers and run simulations at  $m = n = 5, 10, 20, 50, 100, 200, 500, 1000$ , respectively.) For each simulation, we randomly generate 1000 trader groups and compute the expected trading






Figure 8:  $IR(m, n)$ , as a function of  $m$  and  $n$  when agents' reservation prices are beta distributions.

inefficiency ratio as defined in (23). As we can see from Figure 7, which is plotted in a log-log scale, the inefficiency ratios decrease at rate  $1/m$  as  $m$  and  $n$  increase. As the buyers' distribution support shifts from  $[0.5, 1.5]$  to  $[1.5, 2.5]$ , the inefficiency ratio decreases, which is expected because more buyers tend to offer higher bidding prices and the number of successful trading agents increases correspondingly.

In Figure 8, we set the seller's reservation price to be a beta distribution  $g = 1 + B(4, 4)$ <sup>6</sup> and shift the buyers' price distribution from  $f_1 = 0.5 + B(4, 4)$  to  $f_2 = 1 + B(4, 4)$  to  $f_3 = 1.5 + B(4, 4)$ . All the other settings remain the same as in the previous experiment. Two interesting observations are worth mentioning: First, the inefficiency frontier in the case where  $f_3 = 1.5 + B(4, 4)$  is slightly higher than that in the case where  $f_2 = 1 + B(4, 4)$ . The reason for this is that buyers' bidding prices aggregate around two in  $f_3$  and sellers' asking prices aggregate around 1.5 in  $g$ , thus  $b_K - s_L$

<sup>6</sup> $B(4, 4)$  is a symmetric distribution with a unique mode at 0.5 and is supported on  $[0, 1]$ .

is generally larger when  $f = f_3$  than that when  $f = f_2$ , which causes more efficiency loss. Second,  $B(4, 4)$  is not bounded below by a positive number, and thus condition (13) and (14) do not hold. However, as shown in Figure 8, the inefficiency ratio still decreases at rate  $1/m$  when  $m$  and  $n$  are large.

## 6 Conclusions

The efficiency of MDA markets has been a puzzle for many years, especially when trades involve multiple units. This is because agents in an MDA market may have complicated trading strategies available to suit their individual interests, and thus they may not necessarily reveal their truthful intentions, which leads to market inefficiency. We shed light on how this puzzle is solved in one MDA market by analyzing how strong ex-post efficiency is achieved. For this market we designed and implemented an MDA e-market mechanism that is strategy-proof with respect to reservation price, weakly or exactly budget-balanced, asymptotically ex-post efficient and individually rational. Our market mechanism also makes sellers unlikely to under-report their supply volume to drive up the trading price.

If we insist that the market be established based on individuals' free will to trade, strategy-proofness is essential because without truthful revelation, we cannot determine whether a market is efficient or not. Our mechanism, and the consequent analysis of its efficiency makes mathematically concrete a long-speculated, but somewhat vague reason why free markets are indeed efficient: they comprise a large population of players whose individual trading behavior has negligible effect on the overall market.

We implemented a prototype e-market using the RETSINA infrastructure based on our design and

tested its efficiency. Experimental results fit our efficiency theory very well. Another interesting experiment would be testing our theory of strategy-proofness by introducing human traders into our e-market. We plan to do this in future work.

In our MDA market, the trading agents are very simple, because the market mechanism prevents them from manipulating the trading price. Or in other words, the trading agents “understand” the mechanism and hence “know” that pursuing complicate strategies will not benefit them. If some of the assumptions of our MDA market do not hold, it may be possible that “smarter” agents can garner more profit by adopting complicated strategies. In these situations, designing more “intelligent” trading agents is appropriate. One such situation is that a market that allows discriminatory bidding, as we discussed in Section 3. Another one is that of a dynamic market. Our MDA market is a one-shot market, we do not consider the problem where agents who were not included in a trade may come back and resubmit bids. If we introduce timing and dynamics into our MDA market, it will become very similar to the stock market, which seems to defy any exact analysis.

## **Acknowledgment**

The authors are grateful to Professor Stephen Spear of GSIA for his helpful discussion in the early stages of this work, and to two anonymous reviewers for their detailed suggestions and helpful comments that improved the exposition and content of this paper.

## References

- BABAIOFF, M. AND N. NISAN. 2001. Concurrent auction across the supply chain. *In Proceedings of the 3rd ACM conference on electronic commerce*, pp. 1–10.
- BARBERA, S. AND M. JACKSON. 1995. Strategy-proof exchange. *Econometrica*, 63(1):51–87.
- CASON, T. AND D. FRIEDMAN. 1996. Price formation in double auction markets. *Journal of Economic Dynamics and Control*, 20:1307–1337.
- CLARKE, E. 1971. Multipart pricing of public goods. *Public Choice*, pp. 17–33.
- DAS, R., J. HANSON, J. KEPHART AND G. TESAURO. 2001. Agent-human interactions in the continuous double auction. *In Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pp. 1169–1176.
- David, H. 1970. *Order Statistics*. John Wiley & Sons. Inc.
- Fudenberg, D. and J. Tirole. 1991. *Game Theory*. The MIT Press.
- GJERSTAD, S. AND J. DICKHAUT. 1998. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29.
- GODE, D. AND S. SUNDER. 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *The Journal of Political Economy*, 101(1):119–137.
- GROVES, T. 1973. Incentives in teams. *Econometrica*, 41:617–631.

- LEHMANN, B., D. LEHMANN, AND N. NISAN. 2001. Combinatorial auctions with decreasing marginal utilities. *In* Proceedings of the 3rd ACM conference on electronic commerce, pp. 18–28.
- MCAFEE, R. 1992. A dominant strategy double auction. *Journal of Economic Theory*, 56:434–450.
- MYERSON R. AND M. SATTERTHWAITE. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281.
- PEARSON, K. 1902. Note on Francis Galton’s difference problem. *Biometrika*, 1:390–399.
- PLOTT, C. AND P. GRAY. 1990. The multiple unit double auction. *Journal of Economic Behavior and Organization*, 13:245–258.
- RETSINA AFC developers Guide. 2002. Carnegie Mellon University, Intelligent Software Agents Laboratory.
- SANDHOLM, T. 2000. eMdiator: A next generation electronic commerce server. *In* Proceedings of the 4th International Conference on Autonomous Agents(AGENTS), pp.73–96.
- SANDHOLM, T. AND S. SURI. 2001. Market clearability. *In* Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI-01), pp. 1145–1151.
- SATTERTHWAITE, M. AND S. WILLIAMS. 1989. The rate of convergence to efficiency in the buyer’s bid double auction as the market becomes large. *The Review of Economic Studies*, 56(4):477–498.

- SHEHORY, O. AND K. SYCARA. 2000. The Retsina Communicator. *In Proceedings of the 4th International Conference on Autonomous Agents.*
- SYCARA, K., M. PAOLUCCI, M. VAN VELSEN, AND J. GIAMPAPA. 2002. The RETSINA MAS Infrastructure. *Autonomous Agents and Multi-Agent Systems (Forthcoming).*
- SMITH, V. 1982. Microeconomic systems as an experimental science. *American Economic Review*, 72:923–955.
- VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37.
- WELLMAN, M., P. WURMAN, K. O'MALLEY, R. BANGERA, S. LIN, D. REEVES AND W. WALSH. 2001. Designing the market game for a trading agent competition. *IEEE Internet Computing*, March-April, pp.43–51.
- WILLIAMS, S. 1991. Existence and convergence of equilibria in the buyer's bid double auction. *The Review of Economic Studies*, 58:351–374.
- WURMAN, P. , M. WELLMAN, AND W. WALSH. 1998. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. *In Proceedings of the 2nd International Conference on Autonomous Agents*, pp. 301–308.
- YOKOO, M. , Y. SAKURAI AND S. MATSUBARA. 2001. Robust double auction protocol against false-name bids. *In Proceedings of the 21st International Conference on Distributed Computing Systems.*