Experimental Market Mechanism Design for Double Auction*

Masabumi Furuhata[1,2], Laurent Perrussel[2], Jean-Marc Thévenin[2], and Dongmo Zhang[1]

[1] Intelligent Systems Laboratory, University of Western Sydney, Australia
[2] IRIT, Université de Toulouse, France

Abstract. In this paper, we introduce an experimental approach to the design, analysis and implementation of market mechanisms based on double auction. We define a formal market model that specifies the market policies in a double auction market. Based on this model, we introduce a set of criteria for the evaluation of market mechanisms. We design and implement a set of market policies and test them with different experimental settings. The results of experiments provide us a better understanding of the interrelationship among market policies and also show that an experimental approach can greatly improve the efficiency and effectiveness of market mechanism design.

1 Introduction

Auction has been used for many years as the major trading mechanism for financial markets and electronic markets. The existing researches on auction mostly focus on the theoretical aspects of a market mechanism, such as incentive comparability, profit optimization, price formation, and so on [1–4]. From the implementation and market design point of view, "market participants and policy makers would like to know which institution or variant or combination is most efficient, but theoretical and empirical work to date provides little guidance", as Friedman pointed out in [5].

In this paper, we introduce a general approach for the design, analysis, and testing of market mechanisms. The design of a market mechanism involves the development of market policies and evaluation criteria. Different from most existing work in experimental economics, we do not restrict ourselves on specific market policies. Rather, we specify a range of general market policies under the certain trading structure, such as double auction, investigate the properties of market mechanisms with different combinations of market policies. In such a way, we can design a variety of market mechanisms and test them for different purposes.

This paper is organised as follows. Section 2 introduces a market model and specifies the market policies that compose a market mechanism. Section 3 presents a set of evaluation criteria for market mechanism design and testing. Section 4 describes the implementation of market policies. Section 5 presents our experimental results. Finally, we conclude the paper with related work.

* This research was partially supported by the Australian Research Council through Discovery Project DP0988750 and UWS Research Grants Scheme. We thank the anonymous reviewers for their constructive comments.

2 The market model

In this section, we introduce a formal model of market based on double auction market structure. Double auction is a typical market structure in which a set of sellers and buyers trade together for the exchange of certain commodities. Figure 1 illustrates the general structure of double auction markets.

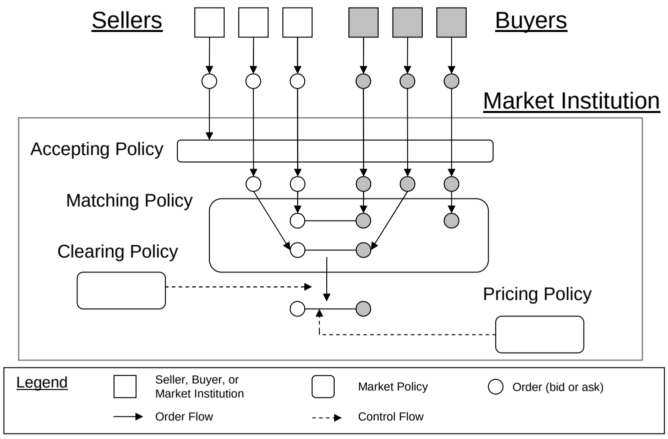

Fig. 1. Structure of Market Institution with Double Auction

In a double auction market, there are three sorts of actors: *sellers*, *buyers* and the *market maker*. The sellers and buyers are called traders. The market maker, who represents the market institution, coordinates the market. During a trading period, the sellers and buyers submit ask orders (sell orders) and bid orders (buy orders) to the market maker of the market institution, respectively. The market maker finds feasible pairs from these incoming orders according to certain market policies, such as accepting policies, matching policies, clearing policies and pricing policies. An *accepting policy* sets criteria for either accepting or rejecting an incoming order. A *matching policy* determines which ask orders match which bid orders. The time for the matched orders to be executed is determined by a *clearing policy*. Meanwhile, the price of the transaction price is determined by a *pricing policy*. According to the structure of double auction markets, the design of market mechanism for a double auction market is to specify each market policy that are to be implemented in the market.

2.1 Market setting

We consider a double auction market of a single commodity. Let $I = S \cup B$ be a set of traders, where S is the set of sellers and B is the set of buyers. We assume that $S \cap B = \emptyset$[3].

Each trader $i \in I$ has a fixed valuation for the commodity, which is private information of the trader, denoted by v_i. Let X be the set of incoming orders. An order $x \in X$

[3] In practice, a trader can be both a seller and a buyer for the same commodity. In such a case, we model it as two different roles because the decision making for selling and buying is different.

consists of two components: the owner of the order, denoted by $\Im(x) \in I$, and the price of the order, denoted by $p(x)$. For any $H \subseteq X$, we write $H^{ask} = \{x \in H : \Im(x) \in S\}$ and $H^{bid} = \{x \in H : \Im(x) \in B\}$. Notice that the meaning of the order prices for sellers and buyers are different. For a seller, the order price (asking price) means that the commodity can be sold with a price no less than this price. For a buyer, the order price (bidding price) means that the commodity can be bought with a price no higher than this price.

2.2 Market Policies

Based on the market setting we describe above, we now define the market policies to govern a double auction market.

An *accepting policy* is a function $\mathfrak{A} : X \to \{1, 0\}$ that assigns to each incoming order a value either 1(accepted) or 0 (rejected). Let $A = \{x \in X : \mathfrak{A}(x) = 1\}$ be the set of all the orders that are accepted under the accepting policy $\mathfrak{A}$.

A *matching policy* is a function $\mathfrak{M} : 2^X \to 2^{X \times X}$ such that for any $H \subseteq X$,

1. if $(x, y) \in \mathfrak{M}(H)$, then $x \in H^{ask}$, $y \in H^{bid}$ and $p(x) \leq p(y)$,
2. if (x_1, y_1) and $(x_2, y_2) \in \mathfrak{M}(H)$, then $x_1 = x_2$ if and only if $y_1 = y_2$.

The first condition sets the feasible condition for a match: *the ask price should be less or equal to the bid price*. The second condition specifies that an order can only be matched once. Let M be a set of matched pairs obtained from the accepting policy, that is, $M = \mathfrak{M}(A)$. We use $|M|$ to denote the number of matched pairs in M.

A *pricing policy* on M is a function $\mathfrak{P} : M \to \Re$ that assigns a positive real number, interpreted as the clearing price, to a pair of ask order and bid order such that for any $(x, y) \in M$, $p(x) \leq \mathfrak{P}(x, y) \leq p(y)$.

Note that any pricing policy is implemented on top of certain accepting policy and matching policy. Without having matched the incoming orders, no transactions can be executed.

A clearing policy determines when to clear a matched pair. Formally, let T be the set of time points of a trading period. A *clearing policy* is a function $\mathfrak{C} : T \times M \to \{1, 0\}$ such that for any $t \in T$ and $(x, y) \in M$, if $\mathfrak{C}(t, (x, y)) = 1$, then $\mathfrak{C}(t', (x, y)) = 1$ whenever $t' \in T$ and $t' > t$, which means that once a matched pair is cleared, it can never come back.

With the implementation of all the above policies, a market maker can determine what, when and how the incoming orders being transacted. Briefly speaking, given a set of incoming orders, the accepting policy determines what orders are to be accepted. Among all accepted orders, the matching policy determines whose good can be sold to whom. The pricing policy determines the actual transaction prices and the clearing policy specifies when the transactions should be executed.

3 Evaluation Criteria of Market Mechanisms

In this section, we propose a set of evaluation criteria for the design and evaluation of market mechanisms. We introduce four indicators to measure profiting efficiency, matching efficiency, transaction volume and converging speed, respectively.

Transaction profit (PR) measures the total revenue of the market institution from all transactions that are executed in a given trading period:

$$PR = \sum_{(x,y)\in M} [c_s(\mathfrak{P}(x,y) - p(x)) + c_b(p(y) - \mathfrak{P}(x,y))]$$

where c_s and c_b is the charging rate of the market institution to a seller and a buyer, respectively. The charging rates represent the percentage of profit that a trader has to pay to the market institution for each transaction.

Allocation efficiency (AE) measures the efficiency of matching policies. Matching a given set of orders is usually referred to as allocation. A set $\hat{M} \subseteq X \times X$ is called a potential matching on X if $(x,y) \in \hat{M}$ implies $x \in X^{ask}$, $y \in X^{bid}$ and $v_{\mathfrak{I}}(x) \leq v_{\mathfrak{I}}(y)$. Let $\mathcal{M}$ be the set of all potential matchings on X. Then the indicator AE measures the rate of the total profit that is made by the current matching policy (resulting the matched set M) against the total surplus between buyers' valuation and sellers' valuation given by the optimal matching on all the incoming orders.

$$AE = \frac{\sum_{(x,y)\in M} (p(y) - p(x))}{\max_{\hat{M}\in\mathcal{M}} \sum_{(x,y)\in\hat{M}} (v_{\mathfrak{I}(y)} - v_{\mathfrak{I}(x)})}$$

Note that the value of the denominator is independent to the currently matching policy while the numerator is determined by the current matching policy. Therefore AE measures the quality of a matching policy.

In many situations, the number of transactions indicates the successfulness of a market. We use the *transaction volume* (TV), i.e., the number of transactions $|M|$, to measure the liquidity of a market.

Finally we use *convergence coefficient* (CC), introduced by Smith [7], to measure the efficiency of a clearing policy. Let

$$CC = \frac{100}{\bar{p}} \sqrt{\frac{(\sum_{(x,y)\in M} (\mathfrak{p}(x,y) - \bar{p}))^2}{n}}$$

where $\bar{p}$ is the average market clearing price. Convergence coefficient is the ratio of standard deviation of transaction prices, which measures the spreads of clearing prices.

4 Implementation of Market Policies

In this section, we briefly describe the approaches we have used for the implementation of each market policy we have introduced in Section 2.2.

4.1 Accepting Policy

An accepting policy defines how incoming orders are accepted by the market maker. A widely used accepting policy is *quote-beating accepting policy* (QBA) under which the market maker accepts an incoming order if it surpasses the current best price among the unmatched orders. Let x_{out}^{ask} and x_{out}^{bid} be the best prices among the current unmatched ask orders and bid orders, respectively. For any incoming order x, the QBA satisfies the following accepting rule:

$$\mathfrak{A}(x) = \begin{cases} 1, & \text{if } \big(\Im(x) \in S \text{ and } p(x) < p(x_{out}^{ask})\big) \text{ or } \big(\Im(x) \in B \text{ and } p(x) > p(x_{out}^{bid})\big) \\ 0, & \text{otherwise.} \end{cases}$$

The QBA accepts an order if it is better than the best prices in the current unmatched orders. This rule is known as New York Stock Exchange (NYSE) rule [8].

The QBA above frequently fails to reduce the fluctuation of clearing prices as pointed by Niu et al. [8]. In order to reduce the fluctuation, they propose *equilibrium-beating accepting* policy (EBA). Let $\tilde{p}$ be the price of the expected competitive equilibrium and δ be an adjustment parameter. For any incoming order x, the EBA makes a binary decision based on the following condition:

$$\mathfrak{A}(x) = \begin{cases} 1, & \text{if } (\Im(x) \in S \text{ and } p(x) \leq \tilde{p} + \delta) \text{ or } (\Im(x) \in B \text{ and } p(x) \geq \tilde{p} - \delta) \\ 0, & \text{otherwise.} \end{cases}$$

Under the EBA, an incoming order is accepted if it exceeds a threshold which consists of the expected competitive equilibrium $\tilde{p}$ and a slack δ. A key issue in the EBA is how to determine the expected equilibrium price $\tilde{p}$ and the slack δ.

We propose a new accepting policy, namely *learning-based accepting policy* (LBA), which requires less parameter tuning than EBA does. A key concept of LBA is to accept an incoming order at higher chances if it is likely to be successfully transacted according to the historical data. Similarly to *linear reward-inaction algorithm* developed by Hilgard and Bower [10], an LBA policy updates the estimation of successful matches at certain prices according to its history data. Let $L : \Re \rightarrow [0, 1]$ be a learning function that assigns an expected matching success rate (a real number between 0 and 1) to an order price. We use two types of learning functions, $L^{ask}(p)$ and $L^{bid}(p)$, for ask and bid orders, respectively. Let $\mathrm{U} = [0, 1]$ be a uniform distribution function. Let $\Pr(\mathrm{U})$ be a probability which is randomly drawn from distribution U. For any incoming order x, an LBA policy determines its acceptance according to the following rule:

$$\mathfrak{A}(x) = \begin{cases} 1, & \text{if } (\Im(x) \in S \text{ and } \Pr(\mathrm{U}) \leq L^{ask}(p(x))) \\ & \quad \text{or } (\Im(x) \in B \text{ and } \Pr(\mathrm{U}) \leq L^{bid}(p(x))) \\ 0, & \text{otherwise,} \end{cases}$$

If a randomly drawn value $\Pr(\mathrm{U})$ is less than or equal to $L(p)$, the market maker accepts an incoming order. A significant part of LBA is how to update the learning function $L(p)$ from the history data. We set the following three types of update rules. Initially, the learning function $L(p)$ is initialised with a constant value α, where $\alpha \in [0, 1]$. The learning function then updates with generated history data. For each successful matched order $(x, y) \in M$, $L(p)$ is increased by a small value ϵ if $0 \leq p \leq p(x)$ or $p(y) \leq p$. For each unmatched order, $x \in X_{out}^{ask}$ or $y \in X_{out}^{bid}$, $L(p)$ is decreased by ϵ if $p(x) \leq p$ or $0 \leq p \leq p(y)$. In other cases, $L(p)$ stays still.

4.2 Matching Policy

Matching policies determine feasible matched pairs between sell orders and buy orders. Since a matching policy is relevant to all the evaluation criteria we have proposed in Section 3, the design of a matching policy is the most important part of market mechanism design. A well-known algorithm is *4-heap algorithm*, proposed by Wurman et al. in [12], which generates efficient and stable matches. The key idea of the 4-heap algorithm is to make matches between the best prices. In other words, the 4-heap algorithm makes matches from the largest bid-ask spreads. In our implementation of market policies, we use 4-heap algorithm for all the experiments we have done.

4.3 Pricing Policy

A pricing policy rules how to set the clearing price for each feasible matched pair. We have implemented three types of well-known pricing policies.

A *mid-point pricing policy* (MPP) is widely used in clearinghouses. The MPP is a unified pricing under which all the matched pairs are transacted at the same price. Among all the matched orders in M, let $x_{in,l}^{ask}$ be the highest matched ask price and $x_{in,l}^{bid}$ the lowest matched bid price. The MPP sets a unified price $\mathfrak{P}(x,y) = \left(p(x_{in,l}^{ask}) + p(x_{in,l}^{bid})\right)/2$ for all $(x,y) \in M$. This pricing is the median of all prices in a given matched set.

In contrast to the MPP, *k-pricing policy* (k-PP) discriminates prices between different matched pairs. Under k-PP, the clearing price is a weighted average of the bid price and the ask price of the pair. The k-PP sets a clearing price $\mathfrak{P}(x,y) = k \cdot p(x) + (1 - k) \cdot p(y)$, where k is a weight parameter ranges in $[0,1]$.

Similarly to k-PP, *N-pricing policy* (N-PP) is a pricing policy, proposed by Niu *et al.* in [8]. Let $M' \subseteq M$ be all the transacted matched pairs. The N-PP determines a clearing price according to the average prices of the matched pairs in M': $\mathfrak{P}(x,y) = \sum_{(x',y') \in M'} \frac{\left(p(x')+p(y')\right)/2}{|M'|}$ for all $(x,y) \in M$.

In addition, we use a pricing policy which always clears at the competitive equilibrium, denoted by CEPP, for benchmarking purpose.

4.4 Clearing Policy

In general, there are two types of double auction markets classified by clearing timing: continuous clearing and periodic clearing. The former clears all matched pairs immediately and the later clears the matched pairs periodically. We implement two types of clearing policies: *continuous clearing* and *round clearing*, where round is the time unit that the traders submit orders in our experiment as we will explain later.

5 Experimental Analysis

We have implemented all the above mentioned market policies on the JCAT platform [14]. The platform has been used as a game server in the Trading Agent Competition in Market Design (TAC-MD) or known as the CAT competition. We have conducted a set of experiments to test the efficiency of each market policy and their combinations.

Our analysis includes identifying the effects of different types of bidding strategies of trading agents. We evaluate the performance of the market mechanisms by considering three types of bidding strategies implemented in the JCAT platform: *zero-intelligence with constraint* (ZI-C), which sets its price randomly, *zero intelligence plus* (ZIP), which is equipped with a learning-based strategy with reactive to market information, and *Roth and Erev* (RE), which is devised by a reinforcement learning strategy to mimic human-like behaviour. We use the same configuration for those bidding strategies as the TAC-MD.

Our experimental settings are described as follows. Each game consists of 20 trading days with several rounds (either 3 or 5 rounds per day) to give learning time for learning-based bidding strategies. Each trader is able to submit a single order in each round; their demand is up to one order in each day. There are 20 sellers and buyers with fixed valuations assigned as follows: $\{\$50, \$55, \ldots, \$145\}$, respectively. Hence, we have the competitive equilibrium at $\$97.5$ and 10 intra-marginal sellers and buyers. Notice that they are the base values for evaluations of market mechanisms. For all experiments, we use the 4-heap algorithm as a matching policy and 10% for the profit fee. For each type of experiments, we run 10 games with the same settings and we evaluate market performances based on the daily average of the evaluation criteria specified in Section 3.

5.1 Pricing Policy Effect

To test market performance under different pricing policies, we fix accept policy to be AAP, matching policy to be 4-heap and clearing policy to be round-based but vary the pricing policies among MPP, k-PP (with $k = 0.5$) and N-PP (with $N = 20$). The experimental results of these three pricing policies are presented in Table 1.

Table 1. Pricing Effects

	ZI-C				ZIP				RE			
Pricing Policy	PR	AE	CC	TV	PR	AE	CC	TV	PR	AE	CC	TV
MPP	15.75	66.08	3.52	4.92	4.64	64.33	15.09	6.01	8.06	34.13	7.90	2.60
k-PP	16.50	67.18	5.21	5.06	4.61	64.79	15.52	6.09	8.16	33.69	7.56	2.56
N-PP	16.50	67.35	1.84	5.06	4.57	63.70	12.72	6.03	7.25	32.38	4.69	2.45

A key result is that pricing policies have a significant impact on the performance of CC but not for the other indicators for all three types of traders. This result suggests us to focus on CC for the design of pricing policies.

A second observation is that N-PP has smaller CC compared to other two policies for all three types of bidding strategies. This can be explained by the volume of historical data required to determine clearing prices. A higher volume helps to improve CC ratio. We use 20 pairs for N-PP, 1 pair for k-PP, and a few pairs for MPP.

Even though ZI-C traders randomly set their bidding prices, the market performance for ZI-C traders is better than the other two market reactive traders. This is particularly true with respect to PR and CC for ZIP traders and PR, AE and CC for RE traders. Therefore, it is interesting whether it is possible to improve these indicators by setting appropriate market policies.

5.2 Accepting Policy Effect

In this experiment, we investigate how accepting policies improve market performance. We compare LBA policy with other three benchmarking accepting policies: AAP, QBA and EBA. We use CEPP for pricing policy, which always clears the market at the middle point of overall traders' valuations. We use continuous clearing for clearing policy. For the LBA policy, we set $\alpha = 1.0$ and dynamically change ϵ in the range between 0.005 and 0.2. We set 3 rounds/day. Table 2 gives the result of our experiments.

Table 2. Accepting Policy Effects

	ZI-C				ZIP				RE			
Accepting Policy	PR	AE	CC	TV	PR	AE	CC	TV	PR	AE	CC	TV
LBA	14.80	91.99	13.23	8.75	11.39	75.00	14.09	6.75	11.06	74.36	14.46	6.74
AAP	13.78	92.14	8.86	8.77	10.80	73.45	10.27	6.72	10.75	72.99	10.34	6.71
QBA	14.76	92.07	8.79	8.78	10.41	72.43	9.77	6.57	10.98	74.59	10.12	6.69
EBA	21.92	84.95	0.44	6.84	15.81	60.15	0.49	4.70	16.10	61.67	0.28	4.80

A key observation is that LBA has better performances than AAP and QBA w.r.t. PR, AE and TV for the market reactive traders (ZIP and RE). This means that LBA signals the market reactive traders in an appropriate way: to accept the orders from the intra-marginal traders and to reject the orders from the extra-marginal traders. This means that LBA can learn the expected matching from history data efficiently. To compare with the results in the previous section, LBA significantly improves AE and TV for all three types of traders. These two indicators are preferable for the traders. However, a disadvantage of LBA is fluctuations of clearing prices as CC indicates.

As a benchmarking accepting policy, we made a set of experiments using EBA policy. We assign the competitive price for the accepting threshold. As a result, the EBA has significant performances on PR and CC. It implies that an EBA policy accepts the intra-marginal traders properly. However, this policy has lower performance on AE and TV, compared to LBA policies. It indicates that there are some unaccepted incoming orders from intra-marginal traders. Hence, the setting of the accepting threshold is too strict for some intra-marginal traders. We detail this point in the next experiment.

In this experiment, we have used a CEPP as pricing policy. In practice, the market institution is not able to obtain the competitive equilibrium. Thus, we consider a case where the market institution uses other pricing policies in the following experiments.

5.3 Robustness of Accepting Policies

Finally, we investigate robustness of accepting policies in a practical setting. We set all the parameters the same as the previous experiment except for the pricing policy. We use N-PP as pricing policy, which has been observed as a well-performing policy in the experiments described in Section 5.1. Notice that N-PP has some fluctuations with clearing prices. Therefore, our aim in this experiment is to investigate the robustness of accepting policies w.r.t. price fluctuation. In addition to the accepting policies used in the previous experiment, we use an EBA policy with slack 5, denoted by EBA(5). The

Table 3. Robustness of Accepting Policies

	ZI-C				ZIP				RE			
Accepting Policy	PR	AE	CC	TV	PR	AE	CC	TV	PR	AE	CC	TV
LBA	15.12	92.67	7.89	8.68	2.42	93.82	2.40	9.36	11.27	74.38	9.22	6.64
AAP	14.40	92.31	8.94	8.89	5.92	84.84	6.60	8.49	10.59	73.38	10.08	6.64
QBA	14.09	91.63	9.08	8.78	2.34	92.27	2.95	9.24	10.49	72.79	10.43	6.67
EBA	22.13	85.66	1.33	6.93	1.24	7.84	4.14	0.70	15.80	61.60	1.45	4.77
EBA(5)	20.15	90.85	1.99	7.91	2.33	90.05	2.30	8.99	15.28	67.58	2.24	5.45

reason for using the slack is to relax the rejection range of incoming orders. The results of the experiment is presented in Table 3.

According to this experimental result, the LBA policy improves all the market performances for all three bidding strategies compared to AAP and QBA policies. Therefore, the LBA policy is more robust for the volatility of clearing prices relative to other compared accepting policies. This may be because a LBA policy is adaptive to the market situations.

An interesting observation is that EBA is no longer high-performing mechanism for ZIP traders if clearing prices are fluctuated. The failure occurs when the clearing price is unbalanced from the competitive equilibrium. In such a case, the strict rejection and the unbalance of the pricing policy make the trader of one-side stop submitting orders, since there is no way to improve for ZIP traders. We also present a case where the rejection range has $\$5$ slack from the competitive equilibrium. The slack improves all the indicators. However, the issues of EBA are how to determine a proper slack and how to estimate the competitive equilibrium.

6 Conclusions and Related Works

In this paper, we have introduced an experimental approach to the design, analysis, and testing of market mechanisms based on double auction. Firstly, we have defined a formal market model that specifies the market policies in a double auction market. Based on this model, we have introduced a set of criteria for the evaluation of market mechanisms. We have designed and implemented a set of specific market policies, including accepting policy, matching policy, clearing policy and pricing policy. A market mechanism is a combination of these market policies. We have conducted a set of experiments with autonomous trading agents to test the efficiency of different market policies.

There are two key findings from our experiments. First, we have observed that a pricing policy has significant effect on the convergence coefficient but does not have similar effect on the other indicators (Section 5.1). This observation suggests that a mechanism designer should focus on the reduction of fluctuations of transaction prices. Second, an LBA policy can help improving allocation efficiency. This is because an LBA policy rejects orders properly if the prices of the orders are not expected to be matched according to the historical data.

The approach that decomposes an auction mechanism into a number of parameterized market policies was introduced by Wurman et al. [15]. In this work, we give a formal definition for each market policies and discussed their properties. Niu et al. [8]

presented a set of experimental results on the market mechanisms with double auction. They proposed N-pricing policy that reduces convergence coefficient and the equilibrium beating accepting policy that increased allocation efficiency. However, their experiments were based on only two simple criteria while our experiments have been based on more comprehensive criteria.

Although the system we have used for the experiments has been developed based on the JCAT platform, we have designed market mechanisms in a single market case in order to focus on the fundamental properties of market policies. This is a major difference from the CAT competition [14] and its analysis [9, 16, 17], since they have dealt with competitive markets. Nevertheless, the findings from the experiments have provided ideas for us to improve market mechanisms for autonomous trading agents. Especially, some of the ideas have been used in the implementation of our TAC Market Design game entry, **jackaroo**, which won the champion of 2009 TAC Market Design (CAT) Tournament.

References

1. Mendelson, H.: Market behavior in clearing house. Econometrica **50** (1982) 1505–1524
2. McAfee, R.P.: A dominant strategy double auction. JET **56** (1992) 434–450
3. Krishna, V.: Auction Theory. Academic Press (2002)
4. Blum, A., Sandholm, A., Zinkevich, M.: Online algorithms for market clearing. J. of the ACM **53**(5) (2006) 845–879
5. Friedman, D.: How trading institutions affect financial market performance: Some laboratory evidence. Economic Inquiry **31**(3) (1993) 410–435
6. Friedman, D.: The double auction institution: A survey. In Friedman, D., Rust, J., eds.: Santa Fe Institute Studies in the Science of Complexity. Addison-Wesley, NY (1993) 3–25
7. Smith, V.L.: An experimental study of competitive market behavior. J. of Political Economy **70**(2) (1962) 111–137
8. Niu, J., Cai, K., Parsons, S., Sklar, E.: Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement. In: AAMAS. (2006) 1143–1150
9. Vytelingum, P., Vetsikas, I.A., Shi, B., Jennings, N.R.: Iamwildcat: Thewinning strategy for the tac market design competition. In: ECAI. (2008) 428–432
10. Hilgard, E.R., Bower, G.H.: Theories of Learning. Prentice-Hall, NJ (1975)
11. Narendra, K.S., Thathachar, M.A.L.: Learning automata-a survey. IEEE Trans. on Systems, Man, and Cybernetics **4**(4) (1974) 323–334
12. Wurman, P.R., Walsh, W.E., Wellman, M.P.: Flexible double auctions for electronic commerce: Theory and implementation. Decision Support System **24**(1) (1998) 17–27
13. Kakade, S.M., Kearns, M., Mansour, Y., Ortiz, L.E.: Competitive algorithms for vwap and limit order trading. In: EC. (2004) 189–198
14. Cai, K., Gerding, E., McBurney, P., Niu, J., Parsons, S., Phelps, S.: Overview of cat: A market design competition (version 2.0). Technical report ulcs-09-005, Department of Computer Science, University of Liverpool, UK. (2009)
15. Wellman, M.P., Walsh, W.E., Wurman, P.R.: A parametrization of the auction design space. Games and Economic Behavior **35** (2001) 304–338
16. Niu, J., Cai, K., McBurney, P.: An analysis of entries in the first tac market design competition. In: IEEE/WIC/ACM Int'l Conf. on IAT. Volume 2. (2008) 431–437
17. Niu, J., Cai, K., Parsons, S., Gerding, E., McBurney, P.: Characterizing effective auction mechanisms: Insights from the 2007 tac market design competition. In: AAMAS. (2008) 1079–1086