

Commodity Trading Using An Agent-Based Iterated Double Auction

Chris Preist

Hewlett Packard Laboratories

Filton Road, Stoke Gifford

Bristol BS12 6QZ UK

Tel: +44 117 922 8311

cwp@hplb.hpl.hp.com

ABSTRACT

This paper describes a new agent-based market mechanism for commodity trading via the Internet. This institution combines the best properties of the continuous double auction and the call auction, but does not suffer from their disadvantages. The institution consists of a marketplace, and a set of agents representing the participants. The agents enter into negotiation with each other in a series of double auctions, and through this determine the equilibrium price of the marketplace. When the equilibrium price has been found, all trades take place at this price.

In the paper, we firstly introduce the concepts of supply and demand, and present the double auction and call auction market institutions. We discuss the advantages and disadvantages of each. We then present the agent-based iterated double auction, and discuss security features that can be incorporated in it. We give one possible implementation of the agent-based iterated double auction, using the PS agents of Preist and van Tol [12], and demonstrate that it can quickly determine the equilibrium price of the market. Finally, we discuss related and further work.

1. INTRODUCTION

This paper describes a new agent-based market mechanism for commodity trading via the Internet.

For many centuries, groups of buyers and sellers have met together in marketplaces to trade with each other. Often, such groups would meet to trade a certain *commodity*, such as wheat or carrots. A commodity is a homogenous good, whose price is the main deciding factor when making a purchase. With the globalisation of trade, international commodity markets for goods such as crude oil and iron ore joined these small informal markets. In these larger marketplaces, more formal rules governing how buyers and seller interact became necessary. Such rules are referred to as a *market institution*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Autonomous Agents '99 Seattle WA USA

Copyright ACM 1999 1-58113-066-x/99/05...\$5.00

The two main market institutions used for trading commodities are the *continuous double auction* and the *call auction*. We will describe these in section 2.

With the advent of the Internet, the creation of global marketplaces has become far cheaper and easier than it once was, since trading can take place by interactions with other participants via the World Wide Web. For example, FastParts [19] has created an international marketplace for trading electronic components, using a double auction. BandX [20] provide a trading forum for telecoms bandwidth. The increasingly open competition within the gas and electricity markets in many countries is creating even more opportunities for such electronic markets.

Agent technology will play an increasingly important role in this revolution. Much work is currently taking place on agents that can negotiate in such marketplaces (For example, [15], [1] and [11]). As the work matures, the speed and continuous vigilance that agents can provide will mean they naturally have an advantage over people in the same marketplace.

However, agents can have another role in electronic trading. We can use agents to design new market institutions – in other words, new rules of engagement between buyers and sellers. Agents can be used to negotiate rapidly and anonymously on behalf of their owners, resulting in frictionless markets that trade at a fair market price and are less open to fraudulent behaviour. In this paper we present the *agent-based iterated double auction*, a new institution which uses agent technology to combine the best features of the continuous double action and the call auction.

2. MARKET INSTITUTIONS FOR TRADING COMMODITIES

We are interested in designing markets that are both efficient and fair. It is appropriate to turn to economics for tools to help us in this. We can use economic principles to inspire the design of agents for use in a market and to analyze the behaviour of the market when they are deployed. Economics is divided into two main sub-areas; microeconomics and macroeconomics. Microeconomics focuses on the structure and dynamics of particular markets, while macroeconomics focuses on the structure and dynamics of entire economies and the effect of government policies on them. As we are interested in studying agents which buy and sell, it is microeconomics which we are concerned with. This section introduces the basic microeconomic

concepts used in this work. We will firstly give an overview of the theory of supply and demand, and then give detailed descriptions of the two main institutions for the trading of commodities; the continuous double auction and the call auction.

2.1 Supply and Demand

Buyers and sellers meet to trade goods and services in a *market*. Buyers may *bid* to buy a good at a given price, and sellers may *offer* to sell a good at a given price. The market has a certain *mechanism*, which determines how bids, offers and other messages can be exchanged to determine a trade. For example, the *English auction* mechanism requires that any buyer may bid for a given good, provided that the bid is higher than the last bid. The bids must be publicly announced. The seller never makes an explicit offer. Rather, they specify a reservation price in advance, and must accept the last bid, provided it exceeds the reservation price. For a review of other examples of auction mechanisms, see [18].

If a good is offered at a given price, the buyers in the market will wish to buy a certain number of these goods. This is said to be the *quantity demanded* at this price. In general, the greater the price of the good, the less the quantity demanded. If we plot the quantity demanded against the price, we get a downward sloping curve, the *demand curve D*. Similarly, if a good is being purchased at a given price, the sellers will be willing to sell a certain number of these goods; the *quantity supplied* at this price. By plotting quantity supplied against price, we can construct the *supply curve S* (See Figure 1). As, in general, sellers will be willing to sell more goods at higher prices, this curve slopes upwards.




Figure 1: Supply and demand

At the price determined by the intersection of the supply and demand curves, the quantity supplied is equal to the quantity demanded. Hence, all participants wishing to make a trade at this price are able to do so. This price is termed the *equilibrium price* P_0 , and the number of goods traded is the *equilibrium quantity*, Q_0 . A *free market* is one in which there is no external intervention, e.g. price controls, and no individual or group of traders who dominate the market. In such a market, trades will naturally tend to take place around the equilibrium price. If

trading is taking place below the equilibrium price, then the quantity demanded is greater than the quantity supplied. There is an *excess demand*. Hence there is an incentive for the buyers to raise their bids to ensure they make a trade. Similarly, if there is an *excess supply*, there is an incentive on sellers to lower their offers to ensure some buyer trades with them. This self-correcting process is known as *price determination* or *equilibration*. The continuous double auction market mechanism described in section 2.1 is particularly effective at rapid equilibration.

Smith [16] introduced a measure of convergence on this equilibrium price, for use in experimental studies on the behaviour of people in double auction marketplaces. This measure, which we will refer to as *Smith's alpha*, is defined as the standard deviation of the actual trade prices from the equilibrium trade price, expressed as a percentage of the equilibrium price:

$$\alpha = \frac{\sqrt{\left(\sum_{i=1}^n (p_i - p_e)^2\right) / n}}{p_e}$$

2.2 The continuous double auction

In the *continuous double auction* (CDA) market mechanism ([7]) buyers and sellers are free at any time to publicly announce bids and offers. Any buyer can accept the offer of a seller, and any seller can accept the bid of a buyer. The CDA originated from informal gatherings of sellers and buyers in local markets, (such as wheat farmers and millers,) and is now a well-established mechanism used in the international financial markets.

One form of continuous double auction that is used for real world trading is the *continuous double auction with order queue* ([17]). In this setup, a trader may make a bid or offer at any time, but once made it persists until the trader chooses to alter it or remove it, or it is accepted. One example of such a marketplace exists on the Internet: FastParts [19] provides a CDA with order queue for buying and selling overstocked electronic components. Buyers and sellers place bids and offers on a web-based trading floor. They revise their bids/offers in response to other trading activity. When a bid and offer meet, they are deleted and a trade takes place. The New York Stock Exchange also uses a form of order queue; the NYSE rule states that the current bid and offer persist, and any new bid or offer must improve on the existing one. However, unlike the FastParts marketplace, a 'reset' occurs when a trade is made, and previous bids/offers must be repeated.

2.3 The Call Auction

The second institution for commodity trading we shall consider is the *call auction*. This differs from the double auction in two major ways:

- There is a central auctioneer who plays an active role in calculating which trades take place.
- All trades take place at the same price.

In the call auction, traders do not publicly announce bids or offers. Instead, they privately prepare information about how many units they would like to buy or sell at a given price.

For example, a seller may decide that if the price is below \$5, they wish to sell no units, if it is between \$5-\$10 they wish to sell 1 unit, and if it is over \$10 they wish to sell 2 units. This information determines the *offer array* for the individual seller. Each seller submits their offer array to the auctioneer, and each buyer submits their *bid array*. The auctioneer combines all the individual offer and bid arrays to determine the total supply and demand in the market. The auctioneer finds the intersection point of these curves, which gives the equilibrium price and quantity.

The auctioneer then announces this price, and matches buyers and sellers who then trade their desired quantity of goods at the equilibrium price. Because all trades take place at the equilibrium price, it is guaranteed that the number of units supplied at this price is equal to the number of units demanded.

2.4 Comparing the institutions

These two institutions each have advantages and disadvantages. The negotiations in the double auction will result in trades taking place at different prices. In the early stages the differences can be quite significant [16]. Prices can start one side of equilibrium and swing to the other side. Traders who could trade at equilibrium may fail to make a trade, while those that can't may succeed, due to this initial instability. The call auction eliminates this element of luck as all trades take place at equilibrium. The call auction is therefore fairer.

Another interesting point of comparison is the time taken to trade. The negotiations in the CDA take time. Because the call auction does not involve negotiation, just calculation, trades are completed more quickly.

However, the call auction has a significant disadvantage; it relies on a central auctioneer. The auctioneer may enter into collusion with some of the participants, and manipulate the market in their favour.

For example, the auctioneer may collude with a seller, and inform them of all the other participants' supply and demand curves. The seller can then determine the best response, and enter that. Often, that will be what they would have entered anyway. However, they may find that the equilibrium price is in a region where there is little competition, and hence increase profits by selling less units at a higher price. (i.e. they may make a monopolistic response)

A *cartel* is a group of traders who manipulate the price by agreeing not to enter into competition with each other. For example, a group of sellers would all agree to sell at the same high price. Cartels are inherently unstable, because there is always an incentive to undercut the competition and make more trades. However, collusion in a call auction can result in cartels being more stable than in a CDA or fair call auction. If a cartel member colludes with the auctioneer, they can check that other members have kept their side of the deal before entering their high-prices supply curve. If they have been betrayed, they can respond appropriately.

The potential for collusion with the auctioneer is a significant disadvantage of the call auction over the continuous double auction. One way around this is to stamp each participant's submission to the auctioneer with a tamper proof timestamp, and

then make the submissions public along with the equilibrium prices. This allows all participants to know that no other participant has swapped a submission and to check the auctioneer's calculation of the equilibrium price. However, this has the disadvantage that the participants supply or demand curves become public after the auction. This is information that many businesses would rather keep secret.

Call auctions either require participants to trust the auctioneer, or to allow potentially sensitive commercial information to be made public. Continuous double auctions, on the other hand, require time consuming negotiations, and result in some trade taking place away from the equilibrium price. By using agent technology, however, we can create a market institution which determines the equilibrium price rapidly and allows all trades to take place at this price. Furthermore, a trusted auctioneer is not needed. The next section describes how this is done.

3. THE AGENT-BASED ITERATED DOUBLE AUCTION

As we explained in the previous section, trades in a continuous double auction often take place at prices away from the equilibrium price. However, experimental work by Smith [16] has shown that if the auction is repeated several times, with participants trading goods with the same values each time, then trades rapidly converge to the equilibrium price as participants respond to market conditions.

This suggests an approach that can allow the double auction to be used to produce trades at equilibrium. Participants engage in a series of mock double auctions and then carry out a final double auction where the trades are actually made. Because the mock auctions have allowed the participants to respond to market conditions, trading should have converged to equilibrium. Hence trades in the final double auction will take place for closer to equilibrium than if the participants had engaged in a single double auction.

Of course, in practice, this approach will not work. There is a strong incentive for participants to attempt to manipulate the market during the mock auctions by refusing to agree trades that in reality they would accept. Because of this, it is unlikely that the series of mock auctions would converge to the actual equilibrium price. Furthermore, this approach would take longer than either the continuous double auction, or the call auction.

However, both of these problems can be overcome if we use agents to trade on behalf of the participants. Cliff [2] has shown that very simple agents can be used in repeated double auctions like those used by Smith, and trading converges rapidly to equilibrium prices. Furthermore Priest and Van Tol [12] have shown that with an appropriate institution this can take place very rapidly in real-time. By using these results we can develop an agent-based version of the iterated double auction described above, which is both immune to attempts at manipulation and performs trades rapidly.

In this institution, a trader will be represented by one or more agents – one agent for each good they wish to trade in this auction. All traders will be issued with identical agents – they simply need to inform each agent if it is to buy or sell and the

highest buying price / lowest selling price they are willing to accept. The agents then enter into the marketplace, and engage in a series of mock double auctions. When the trade prices in these double auctions have stabilised, the marketplace notes the trades and prices agreed in the last double auction, and informs the traders of these. While the double auctions are in progress, the traders can observe the behaviour of their agents in the marketplace, but can't communicate with their agents to adjust the reservation price.

In more detail, the agent-based iterated double auction proceeds as follows:

1. Prior to the auction, all participants receive a copy of the agent. The agent is inspectable, but cannot be altered. Participants can make as many copies as they wish.
2. Participants privately prepare information about how many units they would like to buy or sell at a given price, as in the call auction. For example a seller may wish to sell no units if the price is below \$5, one unit if the price is between \$5 and \$10, and two units if the price is over \$10.
3. Traders then enter appropriate reservation prices into their agents. For example, the seller described above will request one agent to sell one unit for a price of \$5 or more, and a second to sell one unit for a price of more than \$10.
4. The agents enter the marketplace and begin trading. When a buyer and seller agree a trade, they no longer participate in this iteration of the auction, but remain in the marketplace to observe. Once an agent has entered the marketplace it is unable to receive communications from its owner until all iterations of the auction are over.
5. Stage 4 is repeated, with the marketplace measuring the standard deviation of trade prices agreed in each auction. When the standard deviation falls below a previously agreed value (for example, 1% of current trade price), the trade agreed by the agents can be considered binding.
6. The owner of each agent is informed of any trades it has agreed to, and exchange of goods and money takes place.

By allowing agents to negotiate to determine the equilibrium price, we have created a new institution that combines the best properties of the call auction and the continuous double auction. In particular:

- In a CDA, trades can take place at various prices. In the agent-based auction, all trades take place at or near the equilibrium price.
- In the call auction, the activities of the auctioneer are not open to immediate inspection. In the agent-based auction all activity in the marketplace is made public immediately, and so can be verified.
- In the call auction, reservation prices must be made public if the auctioneer's calculations are to be verified. In the agent-based auction, reservation prices remain encapsulated within the agent.
- As we will discuss in section 4, negotiation required in the agent-based auction can take place very quickly, making the agent-based auction as rapid as a call auction.

This is the basic auction protocol. However, if the agent-based auction is to be secure and tamper-proof, we must ensure that it satisfies certain properties.

Firstly, we want all participants to be able to check that everyone is using the correct agent. Without this, it would be possible to enter agents that could adopt a strategy that explicitly exploited the iterated double auction. For example, the agent could use a false reservation price to encourage convergence of the market away from equilibrium, and then swap to their true reservation price when the market had converged. For the same reason, we want to be able to guarantee that no participant is able to communicate with their agent during the iterated double auction, to prevent them from altering the reservation price in the same way. We would also like to ensure that no participant could determine the supply or demand curve of any other participant, either during or after the auction.

One possible solution to these security issues is to modify the protocol described above in two ways. We introduce an anonymizing service between the participants and the marketplace, and we allow any participant to run a mirror of the actual auction to check the outcome. The protocol is amended as follows:

1. Participants determine their reservation prices as described above. However, each participant encrypts the information, and enters the encrypted form in their agents.
2. The agents are not sent directly to the marketplace but are sent firstly to an anonymizing service.
3. As agents are received, they are anonymized and sent on to the marketplace. Any participant can have a secondary marketplace on their system and will also receive copies of the agents. In all cases, the agents are verified by computing their hash value, and ensuring it matches the expected one.
4. When the market is about to begin, participants send their encryption keys to the anonymizing service, which distributes them to the appropriate agents. This allows the agents to decode their reservation price.
5. All marketplaces execute the iterated double auction. Participants verify that their local marketplaces converge to the same equilibrium price as the main marketplace.
6. The results from the central marketplace are sent to the anonymizing service, which forwards them to the participants. Exchange of goods and money takes place.

This additional complexity ensures that all participants can verify that no agent is substituted or altered in the main marketplace during the auction. If such a substitution took place (through collusion between a participant and the marketplace owner) then the local copy of the marketplace would converge on a different equilibrium value. It also ensures that it is very difficult (though not impossible) for a participant to discover the supply or demand data of another player. To do this, they would need to collude with the anonymizing service provider to identify the agents belonging to the participant, and reverse engineer those agents when they arrive in the local marketplace.

4. THE AGENT ALGORITHM AND THE MARKETPLACE

Various designs of the agent and marketplace are possible. However, there are certain criteria they must meet. We would like the behaviour of the marketplace to be straightforward, deterministic and observable. This will allow participants who view the marketplace to ensure that it is behaving fairly. We also want the agents behaviour to be such that the prices of trades converge towards the equilibrium price of the market, to ensure that the iteration terminates at the correct price. Finally, we would like the institution to be rapid, completing all the iterated auctions in as short a time as possible.

Here we present one possible solution using PS agents, which are described more fully in Preist and van Tol [12]. PS agents are based on ZIP agents of Cliff and Bruten [5].

Agents in the community buy and sell an abstract commodity from each other. The agents are divided into buyers and sellers, with each agent wishing to trade one good. Each agent is given its own limit price; if it is a buyer, it will not buy for over this price, **and** if it is a seller, it will not sell for less than this. They are free to make any bid/offer subject to this constraint, and prefer to make a trade at their limit price than to not trade at all. Agents continue trading until all agents have bought/sold or are no longer willing to adjust their bid/offer. At this time, the auction is over. The auction is iterated, with each agent again trying to buy or sell a good for the same limit price.

The limit prices given to the agents determine the underlying supply and demand curves. For example, consider an experiment with 5 buyer agents and 5 seller agents. Let the buyer agents b_1, \dots, b_5 be given limit prices of \$0.50, \$1.00, \$1.50, \$2.00 and \$2.50 respectively. Similarly, let the seller agents s_1, \dots, s_5 also be given limit prices of \$0.50, \$1.00, \$1.50, \$2.00 and \$2.50 respectively. This means that if the good is being traded at \$1.00, then buyers b_2, b_3, b_4 and b_5 each wish to buy one unit in a day, and hence the quantity demanded is 4 units. Similarly, sellers s_1 and s_2 wish to sell, and hence the quantity supplied is 2 units. In this way, we can calculate the quantity supplied and demanded at different prices, and plot the supply and demand curves (figure 2). In this case, the curves intersect at an equilibrium price of \$1.50. At this price, three goods will be traded.




Figure 2: Agents supply and demand

The market institution we use is a form of CDA with order queue. Within each auction, we divide time into discrete rounds. In the first round, all agents participating in the auction must make an

opening bid/offer. On My subsequent round, an agent can update its bid/offer if it so chooses, otherwise its existing bid/offer will stand. If the current highest bid is equal to or greater than the current lowest offer, the marketplace informs the two agents that they have successfully made a trade. The trade takes place at the average of the two prices. Note that the behaviour of the marketplace is deterministic, and can easily be verified by inspection of the trades taking place.

4.1 The Agent Algorithm

Following Cliff [2], our agents consist of a small number of common-sense heuristics combined with a simple learning rule.

Each agent has a profit margin μ , which determines the price at which it is willing to buy or sell relative to its limit price, L . For a seller, μ must lie in the range $[0, \infty)$, and the minimum price p at which it is willing to sell its good is given by $L(1+\mu)$. If the agent makes an offer, it will offer to sell its good at price p . If it receives a bid of a price at or above p , it will accept the bid, and will reject other bids. Similarly, for a buyer μ must lie in the range $[-1, 0]$, and the maximum price p the buyer is willing to pay for the good is given by $L(1+\mu)$. If the agent makes a bid, it will bid to buy a good at price p . It will accept any offer at a price p or below, and will reject other offers. The value p is the *current valuation* the agent places on the good.

Initially, each agent selects a random profit margin in the appropriate range. Each agent then monitors bids, offers and trades in the marketplace, and uses its algorithm to modify its profit margin so as to maximise profit. If it sets its profit margin low, it will not make as much profit as if it sets its profit margin high. However, if it sets its profit margin too high relative to the market, it will fail to make a trade. The agent must use information about current market activity to find the balance, and must respond to changes in the marketplace if a new balance is appropriate.

The agent algorithm runs once in each market round and consists of two phases. Firstly, the heuristics use current market activity to determine what the target profit margin is. Then the learning rule is used to determine how much the profit margin is altered towards the target.

The heuristics we use are simpler than those used by Cliff and Bruten [5]. They are as follows:

Let B_{max} be the highest bid at the beginning of this round, and S_{min} be the lowest offer. Let δ be a small random value. The target value τ for agents to adjust towards are determined as follows:

For BUYERS;

If $S_{min} > B_{max}$ then
target $B_{max} + \delta$

If $S_{min} \leq B_{max}$ then
target $S_{min} - \delta$

For **SELLERS**;

If $S_{min} > B_{max}$ then
 target $S_{min} - \delta$

If $S_{min} \leq B_{max}$ then
 target $B_{max} + \delta$

We also place an additional constraint: If an agent currently has no good to trade, it should not reduce its profit margin. If the above rule requires it to do so, it does not adjust its valuation.

For **results** presented in this paper, we follow Cliff and Bruten [5] in our definition of δ ¹:

If the target is $B_{max} + \delta$
 then $\delta = r_1 B_{max} + r_2$

If the target is $S_{min} - \delta$
 then $\delta = r_1 S_{min} + r_2$

where r_1 and r_2 are independent random variables identically distributed in the range [0,0.2].

The intuition behind these heuristics is straightforward. If trades are not taking place, an agent should attempt to be the most competitive by making the best bid/offer, so should target a valuation slightly better than its competition. If, on the other hand, trades are taking place, an agent should target a valuation slightly better than the best price at which it can currently obtain a trade. Targeting a little better than the current best price allows the agent to 'test' the market, attempting to squeeze a little more profit out.

Given the target value, the agent does not jump straight to that value, but moves towards it at a rate determined by the learning rule. The learning rule used is *Widrow-Hoff with momentum*, which is also used for back propagation learning in neural networks [13]. The learning rule has two parameters. The *learning rate* β determines the speed with which the adjustment takes place, and the *momentum* γ acts to damp oscillation. Given $p(t)$ and $\alpha(t)$, the valuation and target price at time t , the learning rule determines the new valuation, $p(t+1)$, as follows:

$$p(t+1) = \gamma p(t) + (1-\gamma)\beta(\alpha(t) - p(t))$$

4.2 Agent performance

We now summarize experimental results on the performance of PS agents in an iterated double auction. In the results presented here, we use a learning rate of 0.3 and a momentum value of 0.05. For further results, and a comparison with ZIP agents, see Preist and van Tol [12].

¹ Other definitions of δ may be equally valid or better. We believe that it is probably best defined as being relative to B_{max}/S_{min} , with no absolute component r_2 .




Figure 3: Experimental setup

In the market shown in figure 3, there are 11 buyers and 11 sellers given appropriate reservation prices to generate symmetric supply and demand curves. The curves intersect to give an equilibrium price of \$2.00.

Figure 4 gives a time series plot of the price of trades made by PS-agents against auction number. In the first auction, trades are spread over a wide range of values. However, the agents quickly adapt and trades rapidly converge so that all trades take place very close to equilibrium value in later auctions.




Figure 4: Time series of trade prices

To see how effectively PS agents converge on equilibrium in this marketplace, we can calculate the value of Smith's alpha, introduced in section 2, of the trades in a given auction. In figure 5, we give the results of 50 experiments using PS agents carrying out an iterated double auction in this marketplace. The solid line is a plot of the mean value of alpha. The dotted lines either side give the mean \pm one standard deviation.




Figure 5: Mean value of alpha for PS agents

As the graphs show, the PS agents **stabilise** to equilibrium price after 2-3 iterations of the auction, and remain stable with a mean alpha value of 0.4%. In the 2nd auction, alpha is already under 2%, and on the 3rd it is under 1%. These three auctions take place very rapidly, taking on average 48 ± 16 trading rounds to complete all three auctions. Hence PS agents can be used to

implement an agent-based iterated double auction which will rapidly converge to equilibrium.

5. CONCLUSIONS AND FURTHER WORK

In this paper, we have presented a new market institution for the trading of commodities. This institution requires traders to use agents of a certain type to negotiate on their behalf. These agents carry out a series of double auction to determine the equilibrium price of the market, and then transact all trades at a price close to this. The resulting institution is faster and fairer than the standard continuous double auction. Unlike the call auction, it does not require a trusted auctioneer or disclosure of supply/demand information. Hence, it combines the best characteristics of the CDA and call auction in a new institution.

Much work has been carried out on using agents to negotiate prices on behalf of traders (eg [15]. Rust, Miller and Palmer [14] have carried out a double auction tournament to investigate appropriate agent strategies in this institution. However, in all this work, the focus has been on agents belonging to individuals, as opposed to our focus on the provision of restricted functionality agents by a market institution. Easley and Ledyard [6] and Gjerstad and Dickhaut [9] have proposed alternate algorithms for convergence in CDAs, which could be used in place of PS-agents in the agent based iterated double auction. The zero intelligence agents of Gode and Sunder [10], however, are not appropriate, as they do not result in price convergence to equilibrium ([4]). Cliff ([3] has shown that genetic algorithms can be used to determine the optimal parameters of agents to speed up convergence in the marketplace.

In general the agent-based iterated double auction will behave well. However, there are two special circumstances where its behaviour is inadequate:

1. When there is a *price tunnel*; in other words, when the equilibrium price is not a single value, but is a range instead. In this case, the institution will converge on this range, but the stopping criteria would never be satisfied.

Exactly how to handle this case is the subject of further research. A simple, though imperfect, solution is to have a limit on the number of iterations of the auction, and then carry out all trades at the average of the trade prices in the final iteration.

2. Certain supply and demand curves, such as box markets, result in very slow convergence to equilibrium both in markets of agents and of humans ([2]). If convergence is exceptionally slow, it may appear that prices have stabilized when they are not at equilibrium. This would result in a discrepancy between the main marketplace and the participants' local marketplaces, and so giving a false signal that fraud has taken place. Further work is needed on appropriate agent algorithms to speed up convergence of the iterated double auction in such situations, to prevent this occurring.

Despite these limitations, the agent-based iterated double auction can provide a fair, fast and secure way of trading commodities in many situations. It represents a new style of institution, with agents forming part of the institution itself, as opposed to trading through it. Agent based technology has great potential, we

believe, for developing other such institutions, resulting in faster and more effective marketplaces

6. ACKNOWLEDGEMENTS

Thanks to Steve Knight for help, to Janet Bruten and Dave Cliff for comments and to Lin Jones for manuscript preparation.

7. REFERENCES

- [1] Chavez, A. Dreilinger, D. Guttman, R and Maes, P. A Real Life Experiment in Creating an Agent Marketplace. Proceedings of the Second International Conference on the Practical Applications of Agents and Multi Agent Systems, 1997.
- [2] Cliff, D. Minimal Intelligence agents for bargaining behaviours in market-based environments. Technical Report HPL-97-91, Hewlett Packard Labs, 1997.
- [3] Cliff, D. Evolving Parameter Sets for Adaptive Trading Agents in Continuous Double-Auction Markets. Proceedings of the Agents'98 workshop on Artificial Societies and Computational Markets, 1998.
- [4] Cliff, D. and Bruten, J. Zero is not enough: On the lower limits of agent intelligence for continuous double auction markets. Proceedings of the 1998 Symposium on Computation in Economics, Finance, and Engineering: Economic Systems.
- [5] Cliff, D. and Bruten, J. Less than Human: Simple adaptive trading agents for CDA markets. Proceedings of the 1998 Symposium on Computation in Economics, Finance, and Engineering: Economic Systems
- [6] Easley, D. and Ledyard, J. Theories of price formation and exchange in double oral auctions. In [8]
- [7] Friedman, D. The double auction market institution. In [8].
- [8] Friedman, D. and Rust, J. (eds). The Double Auction Market: Institutions, Theories and Evidence. Addison-Wesley 1992.
- [9] Gjerstad, S. and Dickhaut, J. Price formation in double auctions. Games and Economic Behaviour, 22(1), pp1-29, 1998.
- [10] Gode and Sunder Allocative efficiency of markets with zero intelligence traders: Markets as a partial substitute for individual rationality. J. Polit. Econ, 101(1):119-137,1993.
- [11] Guttman, R. Moukas, A. and Maes, P. Agent Mediated Electronic Commerce: A Survey. Knowledge Engineering Review, June 1998.
- [12] Preist, C. and van Tol, M. Adaptive Agents in a Persistent Shout Double Auction. Proceedings of the 1st International Conference on the Internet, Computing and Economics, ACM Press 1998.
- [13] Rumelhart, D. Hinton, G. and Williams, R. Learning internal representations by error propagation. In D. Rumelhart and J. McClelland. Parallel Distributed Processing, Volume 1: Foundations. MIT Press 1986.
- [14] Rust, J. Miller, J. and Palmer, R. (1992) Behaviour of trading automata in a computerized double auction marketplace. In [8].

- [15] Sierra, C. Faratin, P. and Jennings, N. A service oriented negotiation model between autonomous agents. Proc. Modelling Autonomous Agents in a Multi Agent World, 1997.
- [16] Smith, V. An Experimental Study of Competitive Market Behaviour. Journal of Political Economics, 70, 111-137, 1962.
- [17] Smith, V. and Williams, A. An Experimental Comparison of Alternative Rules for Competitive Market Exchange. In Auctions, Bidding and Contracting: Uses and Theory, ed. Engelbrecht-Wiggans. New York University press, pp 307-334, 1983.
- [18] Agorics inc. URL:
<http://www.agorics.com/new.html>
- [19] FastParts URL:
<http://www.fastparts.com/>
- [20] The Bandwidth exchange URL:
<http://www.band-x.com/>