



Adaptive Agents in a Persistent Shout Double Auction

Chris Preist, Maarten van Tol
HP Laboratories Bristol
HPL-2003-242
December 4th, 2003*

E-mail: cwp@hplb.hpl.hp.com, maar@hplb.hpl.hp.com

agents,
double
auction,
electronic
trading,
negotiation,
adaptive
behaviour

Cliff (1997) has demonstrated that simple, adaptive agents are able to trade in a form of double auction marketplace, in such a way that trade prices converge towards the equilibrium price of the marketplace. However, the marketplace within which the agents trade is unrealistic. In this paper, we consider a more realistic form of double auction market, the *persistent shout double auction*. We present agents based on the ZIP agents of Cliff (1997), but with an alternative set of heuristics for use within this auction. We demonstrate that the resulting agents achieve equilibrium significantly faster than ZIP agents do, maintain a more stable equilibrium, and are more robust to changes in learning rate.

1.0 Introduction

In the last few years, there has been an increase in research that combines the disciplines of multi-agent systems and economics, to the mutual benefit of both. This research has made progress in several directions, including;

Market-based control - the use of computerised free market trading to control the allocation of scarce resources. (Clearwater (1996))

Agent-based Computational Economics - The exploration of economic problems by agent-based simulation. (e.g. Tesfatsion (1997))

Electronic commerce - The use of agents to trade on behalf of humans in electronic trading environments. (e.g. Chavez, Dreilinger, Guttman and Maes(1997))

In this paper, we focus specifically on *bargaining agents*; agents which communicate simply by bidding to buy or offering to sell a good at a given price. We present algorithms that can be used by such agents, allowing them to determine a stable market price in a rapid, efficient, robust and distributed fashion. As Cliff(1997) argues, such

algorithms can be applied in all three research areas listed above. In market-based control, bargaining agents can give truly automated and distributed control systems, rather than relying on a central auctioneer (Clearwater, Costanza, Dixon and Schroeder (1996)) or human intervention (Baker (1996)). In agent-based computational economics, bargaining agents can yield insights into how simple human markets converge on equilibrium price, and can place a lower bound on the intelligence necessary for economic activity.

We are particularly interested in the application of bargaining agents in electronic commerce, to trade automatically on behalf of individuals and organisations. Agents are already playing an important role in electronic commerce. Guttman, Moukas and Maes(1998) have shown how agent technology can contribute to different aspects of consumer buying – deciding what to buy, whom to buy it from, how much to pay, and finally the actual exchange of money and goods.

Agents such as BargainFinder [BF] and Jango [JG] are now actively used in e-commerce to search the web and find the trader selling a given product at the cheapest price. Other agents help a customer determine exactly what it is they wish to buy. Firefly [FF] compares a user's taste in music or film with a large database of other people's preferences. It recommends that the user try products that are highly rated by other people who have similar preferences. PersonaLogic [PL] helps a user select the best model of a product, such as a car, for their needs based on a series of questions and answers.

Internet auctions are becoming an increasingly popular way of determining the price at which a good is sold. There is a simple agent, SmartBidder [SB], that will bid on your behalf in an English auction.

However, all these are focussed on business-to-consumer transactions. There are currently no agents on the Internet which focus on business-to-business trading. This is because such transactions require negotiation, and solving the problem of automated

negotiation is not easy. Currently, representatives of the companies involved often negotiate by phone or email, even when the trade is taking place on the Internet. The time and effort to do this means they only negotiate with a small number of possible trading partners, missing good deals elsewhere. However, we expect the world to be very different in the near future. Responsibility for much of this negotiation will be handed over to software agents. These agents will monitor other trade agents continuously, watching for potential opportunities. They will be able to enter into negotiation with many potential trade partners at once, reaching an acceptable deal and setting up a contract in a matter of milliseconds.

If an agent is to negotiate on behalf of an organisation, it needs:

- A representation of the goods or services which are to be traded.
- An understanding of what the organisation wishes to achieve from the negotiation.

- A strategy for negotiation which is at least as effective as a qualified person in the same situation.

Representing the goods to be traded can be complex. It is important to develop a formal representation which makes it possible to specify what a product is, and also how it differs from similar products made by other manufacturers. What features does it have? How well does it perform? This must be done in a way that is perceived as fair by all businesses involved. The goals of an organisation during negotiation can often be very subtle and ill defined. A buyer may be willing to pay more for a higher quality product, or for an added feature. Such tradeoffs are often made instinctively. Hence it is very difficult to capture the exact criteria behind such decisions, allowing an agent to automate them. However there is a class of goods, *homogenous commodities*, for which these two factors are less of an issue.

A good is a homogenous commodity if price is the only factor that is considered when trading it – it cannot be differentiated by being of superior quality, or by adding extra features.

Goods such as crude oil, electricity and wheat are all commodities. Other goods, such as memory chips, network bandwidth and even personal computers, are close to being commodities. Even certain services, such as translation or contract programming, can be treated as commodities. Because price is the only factor involved in comparing one potential deal with another, it becomes easier to represent the goods traded and the organisation's goals. We believe the first deals to be negotiated by automated agents will be for commodity-like goods. For this reason, we will focus on such goods in this paper.

The work presented in this paper shows how simple adaptive agents can negotiate on behalf of buyers and sellers to determine an appropriate trade price in an environment where many buyers are simultaneously negotiating with many sellers. This work is inspired by the ZIP agent work of Cliff (1997) and Cliff & Bruten (1998b). They demonstrate that simple adaptive agents, consisting of a small number of heuristics and a simple learning rule, will learn to trade at equilibrium price in a form of double auction marketplace.

The agents perform in a way that is comparable with humans in a similar experimental setup (Smith (1962)) However, the marketplace they use is unrealistic; at any given time, only one agent can announce a bid/offer. This agent is chosen at random by the market institution. In this paper, we consider a more realistic marketplace, the persistent shout double auction. In a persistent shout double auction, a trader's current bid or offer will persist until the trader makes another. We present an alternative algorithm for agents to use in such a marketplace, and demonstrate that it results in more rapid convergence to equilibrium than the ZIP agents.

In section 2, we present the underlying economic principles used in this work. In section 3, we present the experimental setup. In section 4, we describe the algorithm for use in a persistent shout double auction. In section 5, we present results that compare the performance of our algorithm with the ZIP algorithm of Cliff (1997). Finally, in sections 6 and 7, we discuss how it relates to other work and how we are currently extending it.

2.0 Microeconomic Background

Economics is divided into two main subareas; 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 allowing agents to buy and sell, it is natural to look to microeconomics to analyse the resulting marketplace. For that reason, we present some basic microeconomic concepts here.

2.1 The double auction market

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.

Of particular interest to us in this paper is the *continuous double auction* (CDA) market mechanism. (Friedman (1992)) In a CDA, 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 (such as wheat farmers) with buyers in local markets, 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*, (Smith and Williams (1983)) or *persistent shout double auction*. 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 [FP] provides a persistent shout double auction 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 persistent shout double auction; 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 shouts must be repeated. For a review of other examples of auction mechanisms, see [Agorics].

2.2 Supply and Demand

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 *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.

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




Figure 1: Supply and demand

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 . In a *free market* (One in which there is no external intervention, such as price controls, and no individual or group of traders who dominate the 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 (1962) 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:

$$\alpha = \frac{\sqrt{\left(\sum_{i=1}^n (p_i - P_0)^2\right)/n}}{P_0}$$

3. Experimental Setup

To assess the behaviour of a community of agents in an automated double auction marketplace, it is necessary to perform experiments with a known supply and demand, and hence a known equilibrium price. Following Cliff (1997), we use a setup which is based on that used by Smith(1962) to study the behaviour of humans in double auctions.

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 in a given trading period. 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 trading period is over, and all agents are reinitialised with an intention to buy or sell one good. We use this setup to explore appropriate algorithms for deciding how to adjust the bid/offer an

agent is making in response to the market dynamics.

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

Smith, in his experiments with humans, used a continuous double auction in which agents can shout bids or offers at any time. In our experiments, we have adopted a mechanism, which is a variant of this – the persistent shout double auction, introduced in section 2. We divide time into discrete rounds. In the first round, all agents participating in the auction must make an opening bid/offer¹. On any subsequent round, an agent can update its bid/offer if it so chooses, otherwise its existing bid/offer will stand. Trades take place when bids and offers meet. If the highest bid is higher than the lowest offer, then the trade is made at the average of the two prices. This is close to the mechanism used by Smith, in that it is also a continuous double auction allowing agents to make bids/offers at any time,

¹ This is not a restriction, in that any agent not wishing to make an initial bid/offer, but wishing to participate in the auction, can make a bid of zero, or a ridiculously high offer.

but differs in that in his marketplaces, agents are likely to forget previous bids/offers after a certain time. Hence, there is an informal ‘timeout’ on existing bids and offers.

Cliff uses a more rigid market mechanism than a continuous double auction. Time is divided into rounds. In a given round, one agent with a good to buy/sell is chosen at random to shout its bid/offer. All other agents hear this, and can respond by accepting. If several agents respond, one is chosen at random to complete the trade.

4. The Agent Algorithm

Following Cliff (1997), our agents consist of a small number of common-sense heuristics combined with a simple learning rule. Each agent has a profit margin m which determines the price at which it is willing to buy or sell relative to its limit price, L . For a seller, m must lie in the range $[0, \infty)$, and the minimum price p at which it is willing to sell its good is given by $L(1+m)$. 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 m must lie in the range $[-1, 0]$,

and the maximum price p the buyer is willing to pay for the good is given by $L(1+m)$. 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 is assigned a random profit margin in the appropriate range.² 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 algorithm runs each market round and consists of two phases. Firstly, the heuristics use current market activity to

² A large upper bound (eg MaxInt) is chosen for the seller's profit margin.

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(1997). They are as follows:

Let B_{max} be the highest bid in this round, prior to trades taking place, and S_{min} be the lowest offer. Let δ be a random value, small with respect to B_{max} and S_{min} . 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 = $B_{max} + \delta$

If an agent currently has no good to trade, it continues to observe the market

and adjust its profit margin. However, it does not post a bid or offer. It is subject to an additional constraint: It should not reduce its profit margin. If the above rule requires it to do so, it does not adjust its valuation.

For experiments presented in this paper, we follow Cliff(1997) in our definition of δ :³

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 (Rumelhart, Hinton and Williams (1986)). The learning rule has two parameters. The *learning rate* \mathbf{b} determines the speed with which the adjustment takes place, and the *momentum* \mathbf{g} acts to damp oscillation. Given $p(t)$ and $t(t)$, the valuation and target price at time t , the learning rule determines the new valuation, $p(t+1)$, as follows:

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

5. Experimental Results

We now present experimental results comparing the performance of our

³ 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 . However, for the purposes of comparing our work with that of Cliff(1997) we adopt this definition.

agents in persistent shout double auctions (PS-agents) with that of the ZIP-agents of Cliff. We use the supply and demand curves used by Cliff(1997), and use a learning rate of 0.3 and a momentum value of 0.05.⁴




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 trading period. In the first trading period, trades are spread over a wide range of values, but the agents quickly learn and trades rapidly converge so that

all trades are taking place very close to equilibrium value.




Figure 4: Time series of trade prices

To compare how effectively ZIP and 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 day. In figures 5 and 6, we give the results of 50 experiments using PS and ZIP agents respectively. 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




Figure 6 Mean value of alpha for ZIP agents

⁴ In Cliff(1997), each agent is assigned a random learning rate in the range [0.1,0.5] and a random momentum in the range [0,0.1]. To make comparison between experiments easier, we perform ZIP and PS experiments using the mean value of this random range.

As the graphs show, the PS agents stabilise more quickly than the ZIP agents, and remain consistently more stable. The ZIP agents reach a mean alpha value of just over 1%, while the PS agents are more stable, with a mean alpha value of 0.4%. The PS agents are more stable in the initial trading periods, and achieve stability significantly more quickly than the ZIP agents – on the 2nd trading period, alpha is already under 2%, and on the 3rd it is under 1%. ZIP agents reach a level of just over 1% after 5 trading periods.

The more rapid stabilisation of PS agents is even more pronounced if we consider the number of trading rounds to reach stability. Recall that a trading period continues until all agents have either traded or no longer wish to alter their last bid/offer. Hence, different trading periods can have different numbers of rounds. Because only one agent shouts in each round in the ZIP setup, it means trading periods will tend to take longer. In table 1, we give the mean number of rounds in each of the first 5 periods of trading, both for ZIP and PS agents.

Table 1 : Mean number of rounds in trading period

Trading Period	PS – number of rounds	ZIP – number of rounds
1	26 ± 10	660 ± 472
2	13 ± 3	316 ± 212
3	9 ± 3	253 ± 157
4	8 ± 2	191 ± 109
5	8 ± 2	150 ± 79

As can be seen, the number of rounds in a given trading period decreases as the system stabilises. This is because fewer bids and offers are needed to reach a trade. It is also noticeable that the ZIP setup requires significantly more rounds to complete a trading period than the PS setup. This, combined with the data above, means that the PS agents take, on average, 48 rounds to reach an alpha value of under 1%, while the ZIP agents take on average, 1570 rounds. This means that PS agents converge to a stable value substantially more quickly than ZIP agents do.

PS agents are also more robust to changes in the learning rate used than ZIP agents are. In figures 7 and 8, we compare the performance of PS agents and ZIP agents when the learning rate is set to 0.7, and momentum remains at 0.05. In this case, PS agents behave

slightly less well than previously – they stabilise after 4 trading periods, with a value of alpha of around 1%. ZIP agents behave more erratically, with values of alpha of around 4%.




Figure 7 : Alpha of PS agents with learning rate of 0.7




Figure 8: Alpha of ZIP agents with learning rate of 0.7

In the extreme case, we can consider the behaviour of the heuristics alone (i.e. a learning rate of 1, and a momentum of 0). Figures 9 and 10 give a plot of alpha against trading period. Again, the PS agents are relatively stable, reaching an alpha value of just above 1% after 5 days. ZIP agents however, are not stable – they have a mean alpha value of 6%, with it often ranging up to 10%.




Figure 9: Alpha of PS agents with heuristics only




Figure 10: Alpha of ZIP agents with heuristics only

Hence, PS agents achieve stability significantly more quickly than ZIP agents do, and are more robust to changes in the learning rate used. In Preist and vanTol (1998), we demonstrate that PS agents perform more effectively than ZIP agents for variety of different supply and demand curves, including rapid shifts in supply or demand.⁵ However, there is one case where ZIP agents outperform PS agents; when there are flat supply/demand curves, and an excess of either supply or demand. Like ZIP agents (and humans),

⁵ We are currently investigating the performance of PS and ZIP traders in environments where the supply and demand curves drift with time.

PS agents do converge on the equilibrium price, but only slowly. However, PS agents require more trading periods than ZIP agents do. This may be primarily because ZIP trading periods last many more rounds than PS agents do. Further work is being carried out to determine if other factors are also involved.

6. Related Work

Cliff (1998) and Van Montfort, Bruten and Rothkrantz (1998) have further studied ZIP agents. Van Montfort et al. have demonstrated that ZIP agents can act as arbitrageurs in segmented markets. Cliff has looked at the evolution of appropriate parameter values of the learning rule by using genetic algorithms.

Gode and Sunder (1993) have performed experiments with zero intelligence (ZI) agents, to demonstrate that market discipline, rather than trader behaviour, is partly responsible for the high allocative efficiency and convergence to equilibrium in double auction markets. Cliff and Bruten (1998a) demonstrate that, while market discipline does produce high allocative efficiency,

convergence towards equilibrium of ZI agents was a consequence of the supply and demand curves used in the experiments, and in general does not occur.

Easley and Ledyard (1992) present a theory of price formation in double auctions, which they used to create an agent for the tournament of Rust, Miller and Palmer (1992). Unlike PS and ZIP agents, this theory distinguishes between an agent's current valuation of a good, and its current bid/offer. It uses data on trades in the previous trading period to give the agent an expectation of where trades will take place in the new trading period. Because of this, it does not make predictions about behaviour in the first trading period.

Gjerstad and Dickhaut (1997) propose agents which use data about previous bids, offers and trades (both in the current trading period, and previous trading periods) to determine the probability a bid/offer will be accepted at any given price. Agents can then calculate the expected utility of any bid/offer, and select the highest. This

interpolation and calculation is complex and computationally costly.

Both these approaches are more complex than ZIP and PS agents, and use more historical information. It remains to determine if this added complexity results in the agent making better deals, or whether equally good performance can be gained from a simple learning rule.

7. Further Work

PS agents are able to perform significantly better than ZIP agents because, at any given time, they have access to complete information as to what other agents are willing to pay at that moment. The persistent shout auction is equivalent to all agents shouting what they are currently willing to trade for in each round, while the ZIP setup only allows one agent to shout each round. We have developed an algorithm, which generalises both the ZIP algorithm and the PS algorithm, and allows any number of agents to shout their bid/offer each round. We have used this to show that the speed of stabilization and the level of stability

improve as the percentage of agents shouting increases (Preist and van Tol (1998)).

We have demonstrated that PS agents are able to perform significantly better than ZIP agents in the experiments described above. However, the experimental setup makes certain assumptions about the mechanisms used in the marketplace, and the task performed by the traders.

It assumes that trading is divided into fixed trade periods, in which all participants wish to trade exactly one good. These periods last until all viable trades have been made, at which time all agents are simultaneously reinitialised with an intention to trade. Real marketplaces do not use these fixed trading periods, and traders can gain an intention to trade at any time.

As a period lasts until all agents have either traded or reached their limit price, it assumes that agents have an unbounded time in which to trade. In real markets, traders will usually have a deadline by which they must trade – for

example, to purchase parts used in manufacturing.

The supply and demand in the marketplace remains static for long periods, shifting rapidly to a new stable point. In real markets, supply and demand can drift erratically.

Our current research is aimed at modifying the agents to be able to handle more realistic marketplaces and tasks in each of these three ways.

References

Baker (1996) Metaphor or Reality: A Case Study where Agents Bid with Actual Costs to Schedule a Factory. In Clearwater (1996).

Chavez, Dreilinger, Guttman and Maes(1997) *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.

Clearwater (1996) ed. *Market Based Control*, World Scientific.

Clearwater, Costanza, Dixon and Schroeder (1996) Saving Energy using Market-based Control in Clearwater (1996)

Cliff (1997) Minimal Intelligence agents for bargaining behaviours in market-based environments. Technical Report HPL-97-91, Hewlett Packard Labs.

Cliff (1998) Evolving Parameter Sets for Adaptive Trading Agents in Continuous Double-Auction Markets. To appear in Proceedings of the Agents'98 workshop on Artificial Societies and Computational Markets.

Cliff and Bruten (1998a) Zero is not enough: On the lower limits of agent intelligence for continuous double auction markets. To appear in Proceedings of the Symposium on Computation in Economics, Finance, and Engineering: Economic Systems (1998)

Cliff and Bruten (1998b) Less than Human: Simple adaptive trading agents for CDA markets. To appear in Proceedings of the Symposium on Computation in Economics, Finance,

and Engineering: Economic Systems (1998)

Easley and Ledyard (1992) Theories of price formation and exchange in double oral auctions. In Friedman and Rust (1992)

Friedman (1992) *The double auction market institution*. In Friedman and Rust (1992).

Friedman and Rust (1992) eds. *The Double Auction Market: Institutions, Theories and Evidence*. Addison-Wesley 1992.

Gjerstad and Dickhaut (1997) Price formation in double auctions. To appear in *Games and Economic Behaviour*.

Gode and Sunder (1993) Allocative efficiency of markets with zero intelligence traders: Markets as a partial substitute for individual rationality. *J. Polit. Econ*, 101(1):119-137,1993.

Guttman, Moukas and Maes(1998) *Agent Mediated Electronic Commerce: A Survey*. Knowledge Engineering Review, June 1998.

Preist and van Tol (1998) Generalising ZIP Economic Agents to allow arbitrary numbers of shouts. Hewlett Packard technical report.

Rumelhart, Hinton and Williams (1986) *Learning internal representations by error propagation*. In D. Rumelhart and J. McClelland. *Parallel Distributed Processing, Volume 1: Foundations*. MIT Press 1986.

Rust, Miller and Palmer (1992). behaviour of trading automata in a computerised double auction marketplace. In Friedman and Rust (1992)

Sierra, Faratin and Jennings (1997) *A service oriented negotiation model between autonomous agents*. Proc. Modelling Autonomous Agents in a Multi Agent World, 1997.

Smith (1962) *An Experimental Study of Competitive Market Behaviour*. *Journal of Political Economics*, 70, 111-137.

Smith and Williams (1983) *An Experimental Study of Alternative Rules*

for Competitive Market Exchange in Auctions, Bidding and Contracting: Uses and Theory, edited by R. Engelbrecht-Wiggans, M. Shubik, R. Stark, New York University Press, 1983, pp. 307-334.

Tesfatsion (1997) *How Economists can get Alife*. In *The Economy as an Evolving Complex System II*, Addison Wesley 1997.

[Agorics] Agorics inc. URL:
<http://www.agorics.com/new.html>

[FP] Fastparts URL:
<http://www.fastparts.com/>

[JG] Jango URL:
<http://www.jango.com/>

[FF] Firefly URL:
<http://www.firefly.com/>

[PL] PersonalLogic URL:
<http://www.personalLogic.com/>

[SB] SmartBidder at Numismatists online URL:
<http://www.numismatists.com/>