

Economic Dynamics of Agents in Multiple Auctions

Chris Preist

Hewlett Packard Labs
Filton Rd, Stoke Gifford
Bristol BS34 8QZ, UK
+44 117 312 8311

Chris_Preist@hp.com

Andrew Byde

Hewlett Packard Labs
Filton Rd, Stoke Gifford
Bristol BS34 8QZ, UK
+44 117 312 8764

Andrew_Byde@hp.com

Claudio Bartolini

Hewlett Packard Labs
Filton Rd, Stoke Gifford
Bristol BS34 8QZ, UK
+44 117 312 8505

Claudio_Bartolini@hp.com

ABSTRACT

Over the last few years, electronic auctions have become an increasingly important aspect of e-commerce, both in the business to business and business to consumer domains. As a result of this, it is often possible to find many auctions selling similar goods on the web. However, when an individual is attempting to purchase such a good, they will usually bid in one, or a small number, of such auctions. This results in two forms of inefficiency. Firstly, the individual may pay more for the good than would be expected in an ideal market. Secondly, some sellers may fail to make a sale that could take place in an ideal market.

In this paper, we present an agent that is able to participate in multiple auctions for a given good, placing bids appropriately to secure the cheapest price. We present experiments to show;

- (i) Current auction markets on the web are inefficient, with trades taking place away from equilibrium price, and not all benefit from trade being extracted.
- (ii) Our agent is able to exploit these inefficiencies, resulting in it making higher profits than the simple strategy of bidding in a small number of auctions.
- (iii) As more participants use our agent, the market becomes more efficient. When all participants use the agent, all trades take place close to equilibrium price, and the market approaches ideal behaviour.

Keywords

Auctions. Electronic Commerce. Agents. Negotiation.

1. INTRODUCTION

As a result of the explosion in popularity of e-commerce [1], more and more companies are providing virtual auction sites. Because of this, if you want to purchase a particular good, there are often many auction sites that are offering it. Furthermore, even a single site such as E-Bay may host many auctions for similar goods. If you really want to get the best price, you must monitor all of these auctions using your web browser, and place bids appropriately.

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, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AGENTS'01, May 28-June 1, 2001, Montréal, Quebec, Canada.
Copyright 2001 ACM 1-58113-326-X/01/0005...\$5.00.

Care must be taken to ensure you don't make more than one purchase. If there are a large number of auctions, this can be quite a daunting task, requiring your undivided attention for a period of time. Furthermore, if you wish to purchase more than one item, (as is often the case in B2B trading,) it becomes almost impossible. Because of this, the majority of auction participants focus on a single auction. The use of automated bid technology on sites such as E-Bay encourages this behaviour. A buyer enters a maximum purchase price in a given auction, and the site automatically places bids in that auction on their behalf, locking the participant into that single auction.

This results in an inefficient market. In many cases, the winner of an auction may have been able to get a better price in a different auction, and the losers may have been able to make a purchase elsewhere. In other cases, the seller loses out. If there are few participants in a certain auction, and they value the good at a low price, the good will sell for below market value or may even fail to meet its reservation price. This is despite the fact that losers of other auctions for similar goods would be willing to pay more for it.

In this paper, we describe an agent able to participate in multiple auctions on behalf of a trader, leading to optimal or near-optimal purchase decisions being made. This agent exploits the inefficiencies of the market, locating the auctions that are closing at the cheapest prices and purchasing from there. Furthermore, we present experiments that assess the impact of such an agent on the microeconomic properties of the market. We show firstly that the trader using the agent makes increased profits over the usual strategy of selecting one auction. We also demonstrate that, as more participants adopt the agent-based approach, the market becomes more efficient. The paper is structured as follows. In section 2, we present the algorithm used by the automated agent to participate in multiple auctions. In section 3, we describe experiments showing the effectiveness of the agent, and the effect of the agent on the microeconomic dynamics of the market made up of multiple auctions. In section 4, we describe related work, and in section 5, we present conclusions and future work.

2. THE AGENT ALGORITHM

The agent aims to purchase one or more identical goods on behalf of its user. It can participate in many auctions for this good, and coordinates bids across them to hold the lowest bids. As auctions progress, and it is outbid, it may bid in the same auction or choose to place a bid in a different auction. The algorithm consists of two parts. Firstly, it has a coordination component, which ensures it has the lowest leading bids possible to purchase the appropriate number of goods. Secondly, it has a belief-based learning and

utility analysis component to determine if it should deliberately 'lose' an auction in the hope of doing better in another.

2.1 Definitions

An *auction house* may run one or more *auctions* for a given good. Each auction a_i offers $n(a_i)$ goods for sale. Auctions are assumed to be English auctions in format, with *bidders* placing bids at the price they are currently willing to pay for the good. A bidder may place more than one bid in a given auction. The $n(a_i)$ goods offered in the auction are sold to the bidders making the $n(a_i)$ highest bids, for the price they bid. In case of two equal bids, the item goes to the earliest bidder. Hence the auction is *discriminatory* – some buyers will pay more than others for the same good. Different auctions impose different rules covering how a bid may be entered or retracted. For the purposes of this paper, we assume that a buyer may not retract a bid, and a buyer may enter a bid provided it is at least a certain minimal increment δ above the $n(a_i)$ th highest bid. See Wurman et.al. [2] for a taxonomy of alternative design decisions in auctions.

Our agent participates in many auctions selling similar goods, spread out between many auction houses. It wishes to purchase m goods in these auctions, and is given a valuation of v on each good by its user. To do this, it monitors the set of auctions currently progressing. For each auction a_i , it observes the $n(i)$ highest bids. In other words, it observes the values of the bids which, if the auction terminated immediately, would result in a successful purchase. We refer to these as the currently *active* bids. An active bid b_1 is *lower in precedence* than b_2 if it is lower in value, or if it is the same value but was submitted later. To represent the reservation price r , we assume that the seller initially places $n(i)$ bids of value $r-\delta$, where δ is the minimum bid increment.

2.2 The Coordination Algorithm

Let L be the number of currently active bids that are held by our agent. (Initially, L will be zero.) To ensure it makes m purchases, it needs to make new bids that result in it having an additional $(m-L)$ active bids. As we shall see, this may require it to make more than $(m-L)$ bids, as it may need to outbid itself.

If the agent is to hold j additional active bids in auction a_i , it must place bids that beat the lowest j of the currently active bids placed by competitors. We define the *beatable- j list* for auction a_i to be the ordered set of the lowest j active competitor bids $\{b^1_1, \dots, b^j_1\}$ (where b^j_1 has highest precedence), together with all active agent bids $\{b^1_1, \dots, b^k_1\}$ with precedence lower than b^j_1 . To beat the bids in this list, the agent must place $j+k$ bids of value $b^j_1 + \delta$ where δ is the minimum bidding increment. The *incremental cost* to the agent of placing these bids, if successful, above the cost that it would have incurred in auction a_i previously, is $j \cdot b^j_1 + \delta - \sum\{b^1_1, \dots, b^k_1\}$. The *beatable-0* list of any auction is defined to be the empty set, and has incremental cost of zero. Obviously, an auction for q goods has no *beatable- j* lists for $j > q$.

The agent now constructs potential *bid sets*. A bid set is a set of *beatable- j* lists that satisfies the following criteria;

1. The set contains exactly one *beatable- j* list from each auction.
2. The *beatable- j* lists contain, in total, exactly $(m-L)$ bids made by parties other than our agent. (In other words, the sum of all $J = m-L$)

In other words, each bid set represents one possible way of placing bids to ensure that our agent will gain an additional $(m-L)$ active bids, and therefore will hold exactly m active bids. We define the incremental cost of each of these bid sets to be the sum of the incremental costs of the *beatable- j* lists in it.

The agent must generate the bid set with the lowest incremental cost. In addition, it must avoid generating bid sets that contain a bid equal to or greater than its valuation of the good, v . Various algorithms can be used to do this. We have adopted a depth first strategy through the space of possible bid sets, pruning areas of the search space which are higher cost than the best solution found so far. Full details are presented in [3]

If there is more than one bid set with identically lowest cost, the agent chooses one arbitrarily. If no such bid sets exist, the agent relaxes condition 2 and finds the smallest i such that at least one bid set exists which contains $(m-L-i)$ bids made by parties other than the agent. Given this i , the agent chooses the bid set with the lowest incremental cost.

Having generated the bid set with the lowest cost, the agent places bids in each auction. For each *beatable- j* list in the bid set, the agent places $j+k$ bids of value $b^j_1 + \delta$ in the corresponding auction a_i .

The agent continues to monitor the auction, and repeats its analysis if other parties place new bids. In this way, the agent ensures it maintains m active bids at the least possible cost to itself, unless doing so requires it to place bids above its valuation of the good. Providing all auctions terminate simultaneously, this will result in it buying the goods at the best price possible, given the competition in each auction.

2.3 Auctions Terminating at Different Times

Now, we consider the case where auctions terminate at different times. In such a situation, the algorithm above will not necessarily behave optimally. Imagine a situation where an auction starts every half-hour, and lasts for an hour. The agent would always monitor two auctions, one that is nearer closing than the other. Inevitably, bids will be higher in the auction that is nearing completion. Hence the agent would switch bidding to the newer auction, and withdraw from the auction about to close. If this continued, the agent would never make a purchase, but would simply switch bids to a new auction every half-hour.

The agent needs a mechanism for determining whether to remain in an auction which is about to close, even when there are other auctions with lower current bid prices. To do this, it must be able to make a trade-off in terms of expected value between the relative certainty of remaining in an auction about to close, against the risk of participating in a newer auction. The newer auction may result in a lower purchase, or may result in a far higher purchase price above the agent's valuation of the good. In this section, we propose a mechanism for doing this.

The mechanism we use combines simple learning with utility theory. The agent uses learning to build a model of the spread of valuations held by participants in different auction houses. Then, based on its beliefs about these valuations, it calculates the utility of likely participation in persisting auctions, and compares this with the certain outcome in the terminating auction. If the terminating auction has a higher utility, it remains a participant and makes the purchase. If the remaining auctions have higher

expected utility, it withdraws from the terminating auction and continues participation elsewhere.

2.3.1 The Learning Mechanism

The agent generates a model of the potential outcome of auctions by creating a model of each auction house. For a given auction house and a given type of good, it creates a belief function $B(x,q)$ representing the probability that x bidders value the good with a valuation greater than q in a given auction for that good. It builds up this function by monitoring auctions for the good conducted by the auction house. Various possible learning techniques can be used to generate this function, and are discussed in [3]. The exact choice will depend on the underlying dynamics of the demand for the good under consideration.

Using this function, we can estimate the probability that a bid of a certain value will be successful in an auction by a given auction house. Consider an auction for n goods, in which our agent wishes to purchase one. The probability that a bid of q by our agent will be successful can be estimated to be $1 - B(n,q)$; i.e. 1 minus the probability that n other bidders are prepared to outbid our agent.

There is a flaw in this model, which must be taken into account if it is to be successful. Unlike a Vickrey or Dutch auction, an English auction reveals nothing about the valuations of successful bidders. In other words, if a bidder makes a successful bid of x , we cannot be sure how much higher they may have been willing to bid. To take account of this, it is necessary to add some kind of heuristic weighting to the belief function – we must increase the value of a successful bid by a certain amount, to reflect this possible willingness to bid higher. One possibility is to add a small random amount to each successful bid. In some domains, it may be possible to use econometric data to determine accurately the range that this should be drawn over, while in other domains it may be necessary to use a heuristic estimate.

2.3.2 Utility analysis of leaving an auction

We now consider how this belief function can be used to compare the expected payoff of an auction that is about to terminate with the less certain outcome of other auctions that terminate later. For the sake of clarity and brevity, we present the technique assuming our agent wishes to purchase a single good.

The expected payoff from the terminating auction is simple to calculate. Assuming our agent is holding an active bid q , or is able to place one at the last moment, then the payoff will be $(v-q)$. If the agent is unable to place a bid because all active bids are beyond its valuation of the good, then payoff will be zero and the agent is forced to participate in other auctions.

The expected payoff of continuing to participate in the non-terminating auctions is more complex to calculate. To do this, we use the belief function to calculate the probability our agent will be able to make a purchase at various possible bid prices. Recall that, for a given bid price q , the probability our agent will make a successful bid in an auction run by a given auction house is $1 - B(n,q)$, where n is the number of goods being sold. Similarly, the probability that our agent will be able to make a successful bid at a lower price, $q-1$, is $1 - B(n,q-1)$. Hence, the probability that our agent will succeed with a bid of q and no lower is $B(n,q-1) - B(n,q)$. The utility of this outcome will be $(v-q)$. Hence, we can calculate the expected utility of participating in a given auction as;

$$\sum_{q=0}^v [B(n,q-1) - B(n,q)](v-q)$$

Of course, as the auction may already be in progress, it is necessary to take into account the current active bids in that auction. The general belief function $B(x,q)$ for the auction house is therefore adapted for this particular auction a_n to give $B(a_n,x,q)$. If the good being traded is a private value good, and hence all buyers have valuations independent of each other, this is defined as follows;

Let p be the value of the x th highest bid in auction a_n

Then $B(a_n,x,q) = B(x,q)/B(x,p)$ for all $q \geq p$

1 for all $q < p$

Given an expected utility on the remaining auctions, the agent must decide whether to place higher bids in the auction that is about to terminate, or withdraw from it. If we assume that the agent is risk neutral, then it will be willing to bid up to a value where the actual utility of the terminating auction is the same as the highest expected utility among the remaining auctions. In other words, it is prepared to make a maximum bid b_{max} of;

$$b_{max} = v - \sum_{q=0}^v [B(n,q-1) - B(n,q)](v-q)$$

In this way, the agent is able to make informed decisions about whether to continue bidding in an auction or to switch. If it is making multiple purchases, it may purchase some in the terminating auction, and choose to switch others to continuing auctions. Extensions of the algorithm to handle this case will be dealt with in a future paper.

3. EXPERIMENTAL ANALYSIS


We now present an analysis of the efficiency of the agents, and their effect on the microeconomic properties of the market.

3.1 Infrastructure

A custom written simulator was used for all experiments. The simulator was written in Java, and ran entirely on one machine. A separate thread handled each auction's operation, and hence auctions operated asynchronously. In a given auction, each trading period was divided into rounds, in each of which there were two steps. In the first step the auction house asked each participating agent for its choice of bids, and accepted those bids that met the improvement criterion, which was that a bid shall be at least 1 higher than the current lowest bid. Duplicate bids are allowed, with earliest posted being considered higher. In the second step, all successful bids shouted in step 1 were broadcast to all participating agents. A run of an experiment ended when no auction received any new bids, at which point goods are awarded to the highest bids, at the bid prices.

3.2 Experimental design

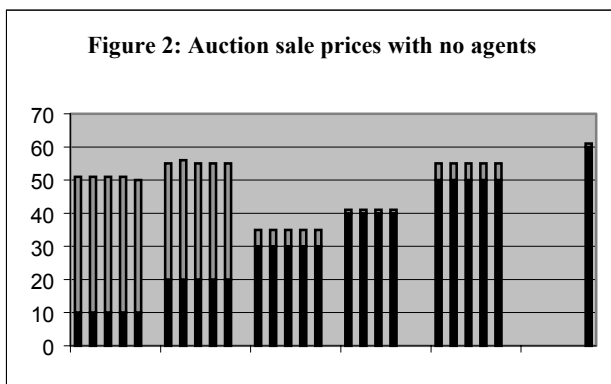
In each experiment there were 6 simultaneous auctions, with reservation prices $\{10, 20, 30, 40, 50, 60\}$. Each auction had 5 goods available for sale. There were always 40 buyers, whose reservation prices were 2 each of $\{35, 40, 45, 50, 55, 60, 65, 70\}$ and the rest with reservation price 30. Each buyer's goal was to buy 3 goods. The utility of a trade to the auction was simply the trade price minus the auction's reservation price. The utility of a trade to the buyer was the buyer's reservation price minus the



trade price, and buyers were constrained to buy as many goods up to a maximum of 3, provided the trade had positive utility.

The resulting supply and demand curves for the market are shown in Figure 1. The ascending supply curve shows that, as price increases, more sellers are willing to sell. The descending demand curve shows that, as prices increase, less buyers are willing to buy. At the intersection point, the number of buyers and sellers wishing to trade is equal. This is the *equilibrium price*, P_0 , and the quantity traded at this price is the *equilibrium quantity*, Q_0 . In this particular setup, the equilibrium price is 50, and the equilibrium quantity is 25. In an ideal market, trade takes place at the equilibrium price, and according to the first theorem of welfare economics, the outcome is pareto optimal. (i.e. all gains from trade are realised.)

The experimental economist Vernon Smith used two measures in his comparison of actual markets with the ideal [4]. *Allocative Efficiency* is defined as the total actual profit earned by all the traders divided by the maximum total profit that could have been earned in an ideal market, expressed as a percentage. This is a measure of how much of the potential gain from trade has been realised within the market. *Smith's Alpha*, α , is a coefficient of convergence which measures how close actual trade prices are to the equilibrium. This is defined as $\alpha = 100\sigma_0/P_0$, where σ_0 is the standard deviation of trade prices around P_0 rather than around the mean price.



The simulation uses two types of bidding algorithm. We represent human bidders with a simple automated strategy. Before the auctions open, the buyer randomly chooses three different auctions, and attempts to purchase one good from each. It then bids against other bidders, up to its reservation price. Automated agents use the algorithm described in section 2. We present a series of experiments, exploring the effect of adding agents to the market. Initially, we run the market with no automated agents, and observe its efficiency. We then allow one buyer to use the agent strategy, and measure its utility gain in comparison with the naïve strategy, for different reservation prices. We then randomly assign automated agents to 0, 1, 5, 10, 20, 30 and 40 of the buyers. In each case we observe the allocative efficiency and Smith's alpha, to identify how close to an ideal market the behaviour is.

3.3 Results

Figure 2 plots the trading prices for the 6 auctions in a single run with no automated agents. Note the spread of trade prices, resulting in a high value of alpha (17.1). In auctions 2, 5 and 6, buyers have paid more than they would have in an ideal market. In auction 4, the seller has failed to make a sale, while in an ideal market, buyers would have been willing to trade at this price. Hence allocative efficiency is low (92%).

Table 1. Utility extracted with and without an agent

Number of other agents	No agent, reservatio n = 45	Agent, reservatio n = 45	No agent, reservatio n = 55	Agent, reservatio n = 55	No agent, reservatio n = 65	Agent, reservatio n = 65
0	1.32	1.44	5.85	9.94	14.5	20.0
9	0.14	0.08	4.89	5.52	13.1	15.8
19	0	0	3.92	4.67	11.6	14.6
29	0	0	4.06	4.36	12.85	14.4
39	0	0	3.25	4.13	12.3	14.2

Table 1 shows the utility gain that can be made by a single buyer adopting the agent approach. It compares the utility gain for a buyer with different valuations, either using the agent or not. The buyer competes against 39 other buyers, with the leftmost column specifying how many of these use the automated agent. The results are the average utility of the agent over 20 or more runs, where the other agents are randomly allocated to buyers. Looking at the first row, we see that the agent based approach yields higher utility than the standard strategy in all cases (Except in the case of 9 agents, with our agent having a reservation price of 45. We believe this is a statistical blip.) In the case of traders with reservation higher than the equilibrium price, the difference is substantial. However, looking down each column, we see that as more traders adopt the agent-based approach, the utility gained by each decreases. When more than half the community adopts the agent-based approach, each buyer makes less profit than they did when the community as a whole didn't use agent technology.

Figure 3a. Relationship between number of autonomous agents, and Smith's alpha

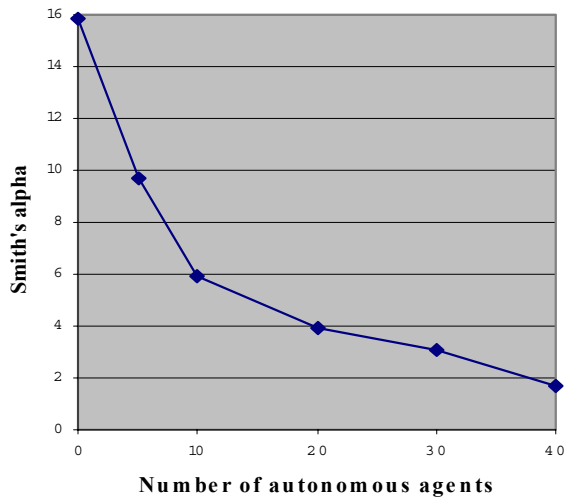
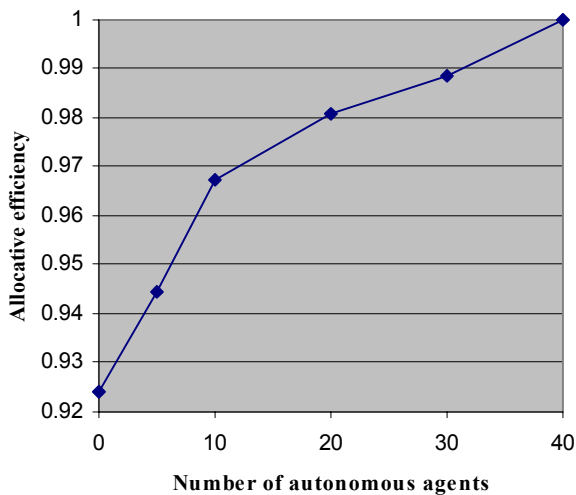
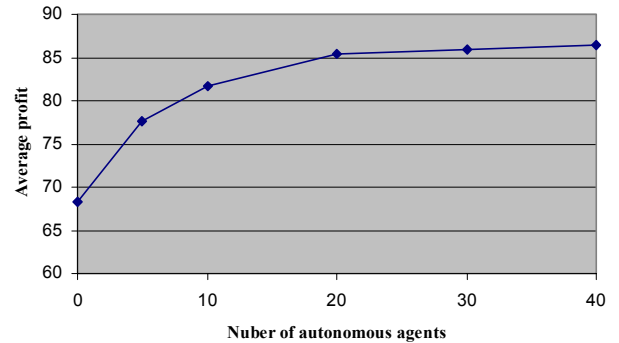


Figure 3b. Relationship between number of autonomous agents, and allocative efficiency



Hence, we have an interesting phenomenon, analogous to a multi-party form of the prisoner's dilemma; It is individually rational for any one buyer to adopt the automated agent, but this sets in motion a trend which is detrimental to each individual buyer, and hence to the social welfare of the community of buyers as a whole. Figures 3a and 3b show how, as more automated agents are added to the market, it becomes more efficient. The curves are the average value of alpha and allocative efficiency over 50 runs. It is interesting to observe that even a small number of agents can result in significant efficiency improvements in the market. This is even more pronounced if, rather than allocating agents randomly to buyers, the agents are adopted by buyers with high valuations.

Figure 4. Relationship between average auction profits and number of autonomous agents



In light of the lower utility that can be extracted if all buyers use an agent, one must ask where all the utility that this higher efficiency corresponds to is going. The answer in the case of English auctions such as the one we study, is to the sellers, as figure 4 shows. It plots the average profit made by each auctioneer, as more buyers adopt the agent. As can be seen, the average profit increases as the market becomes more efficient. The submarginal auctioneers (those with reservation prices higher than equilibrium price) lose out, failing to trade at all. However, this is more than compensated for by the additional profits made by the others. Hence, it is in the interests of auctioneers (unless they are submarginal) to promote the use of agents in their auction houses, as opposed to the simple automated bidding strategies offered currently.

4. Related Work

Research into automated negotiation has long been an important part of distributed AI and multi-agent systems. Initially it focused primarily on negotiation in collaborative problem solving, as a means towards improving coordination of multiple agents working together on a common task. Laasri, Lassri, Lander and Lesser [5] provide an overview of the pioneering work in this area. As electronic commerce became increasingly important, the work expanded to encompass situations with agents representing individuals or businesses with potentially conflicting interests. The contract net [6] provides an early architecture for the distribution of contracts and subcontracts to suppliers. It uses a form of distributed request-for-proposals. However, it does not discuss algorithms for determining what price to ask in a proposal. Jennings et.al. [7] use a more sophisticated negotiation protocol to allow the subcontracting of aspects of a business process to third parties. This is primarily treated as a one-to-one negotiation problem, and various heuristic algorithms for negotiation in this context are discussed in [8]. Vulkan and Jennings [9] recast the problem as a one-to-many negotiation, and provide an appropriate negotiation protocol to handle this. Other relevant work in one-to-one negotiation includes the game-theoretic approach of [10] and the logic-based argumentation approach of [11].

As much electronic commerce involves one-to-many or many-to-many negotiation, the work in the agent community has

broadened to explore these cases too. The Michigan AuctionBot [12] provides an automated auction house for experimentation with bidding algorithms. The Spanish Fishmarket [13] provides a sophisticated platform and problem specifications for comparison of different bidding strategies in a Dutch auction, where a variety of lots are offered sequentially. The Kasbah system [14] featured agents involved in many-to-many negotiations to make purchases on behalf of their users. However, the algorithm used by the agents (a simple version of those in [8]) was more appropriate in one-to-one negotiation, and so gave rise to some counter-intuitive behaviours by the agents. [15] and [16] present adaptive agents able to effectively bid in many-to-many marketplaces, and are the first examples of work which borrow techniques from experimental economics to analyze the dynamics of agent-based systems. [17] demonstrates how these can be used to produce a market mechanism with desirable properties. Park et.al. [18][19] present a stochastic-based algorithm for use in the University of Michigan Digital Library, another many-to-many market.

Gjerstad et. al. [20] use a belief-based modeling approach to generating appropriate bids in a double auction. Their work is close in spirit to ours, in that it combines belief-based learning of individual agents bidding strategies with utility analysis. However, it is applied to a single double auction marketplace, and does not allow agents to bid in a variety of auctions. Vulkan et.al. [21] use a more sophisticated learning mechanism that combines belief-based learning with reinforcement learning. Again, the context for this is a single double auction marketplace. Unlike Gjerstad's approach, this focuses on learning the distribution of the equilibrium price. Finally, the work of Garcia et.al. [22] is clearly relevant. They consider the development of bidding strategies in the context of the Spanish fishmarket tournament. Agents compete in a sequence of Dutch auctions, and use a combination of utility modeling and fuzzy heuristics to generate their bidding strategy. Their work focuses on Dutch rather than English auctions, and on a sequence of auctions run by a single auction house rather than parallel auctions run by multiple auction houses. However, the insights they have developed may be applicable in our domain also. We hope to investigate this further in the future.

In this paper, we have shown how agents operating in multiple auctions can create a more efficient market. An alternative approach is to attempt to provide the right market mechanism in the first place, providing a centralized point of contact for all buyers and sellers to trade. Sandholm [23] proposes a sophisticated marketplace able to handle combinatorial bidding, and able to provide guidance to buyers and sellers as to which market mechanism to adopt for a particular negotiation. In the long term, as the different auction houses merge or fold and only a few remain, this approach will be ideal. In the short term, we expect improved market dynamics will occur through autonomous agents in multiple auctions.

It is interesting to contrast our analysis with that of Greenwald and Kephart [24]. They demonstrate that the use of dynamic price-setting agents by sellers, to adjust their price in response to other sellers, can lead to an unstable market with cyclical price wars occurring. We, however, show that (in a very different context) the use of agents improves the dynamics and stability of the market. From this, we can conclude that agent technology is not a-priori 'good' or 'bad' for market dynamics, but that each potential role must be studied to determine its appropriateness.

5. CONCLUSIONS AND FUTURE WORK

We have presented an agent able to participate simultaneously in multiple auctions, and explored its properties. The experimental analysis has demonstrated that the agent strategy outperforms the standard single auction strategy significantly, and also leads to more efficient markets. However, as more participants adopt the agent-based approach, the community of buyers lose out, and end up making less profit than they made when no-one was using an agent. The auctioneers (sellers) are the ones who benefit most from the extra efficiency of the market.

In this paper we have presented experiments using a single, specific supply/demand graph. We intend to carry out further experiments with different supply/demand curves, and mathematical analysis, to determine if the sellers receive the additional surplus in all circumstances, or if the buyers may find the use of agents to be collectively as well as individually beneficial.

The experiments presented in this paper focus on auctions which run in parallel, and hence the automated agent does not need to make use of the utility mechanism to select between current and future auctions. Additional experiments will be carried out, exploring the effect of staggering and randomizing the auction start and end times, to explore the effect of this on the market dynamics and the capabilities of the agent to handle this.

In addition, we plan to extend the functionality of the agent. We will explore how effective different learning mechanisms are, and also adapt the agent to be able to participate simultaneously in a variety of different auction environments – eg Dutch, Vickrey, Continuous Double Auction, etc.

6. ACKNOWLEDGEMENTS

Thanks to Ian Dickinson for development of the experimental testbed, to Ivan Phillips for initial prototyping, and to Kenny Paterson for mathematical support.

7. REFERENCES

- [1] Kalakota, R. and Whinston, A. *Frontiers of Electronic Commerce*. Addison-Wesley 1996.
- [2] Wurman, P.R. Wellman, M.P. and Walsh, W.E. A parameterization of the auction design space. *Games and Economic Behavior*, to appear, 2000.
- [3] Preist, C.W, Bartolini, C. and Phillips, I. *Algorithm Design for Agents which Participate in Multiple Simultaneous Auctions*. In *Agent Mediated Electronic Commerce III (LNAI)*. Springer-Verlag: Berlin, Germany. To appear, 2000.
- [4] Smith, V.L *Papers in Experimental Economics*. Cambridge University Press, Cambridge 1992
- [5] Laasri, B. Laasri, H. Lander, S. and Lesser, V. A Generic Model for Intelligent Negotiating Agents. *International Journal of Intelligent and Cooperative Information Systems*, 1(2) 1992 pp291-317.
- [6] Smith R.G.. The contract net protocol: high-level communication and control in a distributed problem solver.. *IEEE Trans. Comput.*, 29, 1104-1113, 1980.
- [7] Jennings, N.R. Faratin, P. Johnson, M.J. O'Brien P.O. and Wiegand, M.E. *Using Intelligent Agents to Manage Business*

- Processes. In Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM-96), 345-360, April 1996
- [8] Faratin, P. Sierra, C. and Jennings, N. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24(3-4),1998, pp159-182.
- [9] Vulkan, N. and Jennings, N. Efficient Mechanisms for the Supply of Services in Multi-Agent Environments, Proceedings of the 1st International Conference on the Internet, Computing and Economics, ACM Press 1998.
- [10] Rosenschein, J. and Zlotkin, G. *Rules of Encounter*. MIT Press, 1994.
- [11] Parsons, S. Sierra, C. and Jennings, N. Agents that reason and negotiate by arguing.
- [12] Wurman, P. Wellman, M. and Walsh, W. The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents. In Proc. Second Conference on Autonomous Agents, 1998.
- [13] Rodriguez-Aguilar, J. Noriega, P. Sierra, C. and Padget, J. Fm96.5: A Java-based electronic auction house. Proc. Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Systems, 1997, pp207-224.
- [14] Chavez, A. Dreilinger, D. Guttman, R. and Maes, P. A Real-Life Experiment in creating an Agent Marketplace. Proc. Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Systems, 1997.
- [15] 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
- [16] 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.
- [17] Preist, C. Commodity Trading using an Agent-Based iterated Double Auction. Proc. Third International Conference on Autonomous Agents, 1999.
- [18] Park, S. Durfee, E. and Birmingham, W. Emergent Properties of a Market-Based Digital Library with Strategic Agents. Proc. International Conference on Multi Agent Systems, 1998.
- [19] Park, S. Durfee, E. and Birmingham, W. An Adaptive Agent Bidding Strategy based on Stochastic Modelling. Proc. Third Conference on Autonomous Agents, 1999.
- [20] Gjerstad, S. and Dickhaut, J. Price formation in double auctions. *Games and Economic Behaviour*, 22(1), pp1-29, 1998.
- [21] Vulkan, N. and Preist, C. Automated Trading in Agents-based Markets for Communication Bandwidth. Proc. UKMAS (1999).
- [22] Garcia, P. Gimenez, E. Godo, L. and Rodriguez-Aguilar, J. Possibilistic-based design of bidding strategies in Electronic Auctions. Proc. 13th Biennial European Conference on Artificial Intelligence, 1998.
- [23] Sandholm, T. eMediator: A Next Generation Electronic Commerce Server. Proc. Fourth International Conference on Autonomous Agents, 2000.
- [24] Greenwald, A.R. and Kephart, J.O. Shopbots and Pricebots. Proceedings of the 16th International Joint Conference on Artificial Intelligence, IJCAI 1999.