An Agent Model for First Price and Second Price Private Value Auctions

A. J. Bagnall and I. Toft

School of Information Systems
University of East Anglia
Norwich, NR47TJ
England
ajb@sys.uea.ac.uk

Abstract. The aim of this research is to develop an adaptive agent based model of auction scenarios commonly used in auction theory to help understand how competitors in auctions reach equilibria strategies through the process of learning from experience. This paper describes the the private value model of auctions commonly used in auction theory and experimentation and the initial reinforcement learning architecture of the adaptive agent competing in auctions against opponents following a known optimal strategy. Three sets of experiments are conducted: the first establishes the learning scheme can learn optimal behaviour in ideal conditions; the second shows that the simplest approach to dealing with situations of uncertainty does not lead to optimal behaviour; the third demonstrates that using the information assumed common to all in private value model allows the agent to learn the optimal strategy.

1 Introduction

The internet has meant that auctions have become a crucial means of conducting business, in both the consumer to consumer market (through such companies as Ebay) and the business to business arena through companies such as Freight-Traders Ltd [2] and FreeMarkets Inc [1]. Research indicates that the B2B market will escalate from $43 Billion in 1998 to $1 Trillion in 2003 [12], and much of this business will be conducted through auctions, which are rapidly replacing request for quotes as the preferred means of business procurement. This growth in on-line auctions has created an increasing interest in the study of the behaviour of bidders and the effect of auction structure on the choice of bidding strategy [3, 5].

Broadly speaking, some of the key questions in studying auctions are: what are the optimal strategies for a given auction structure; how do agents reach an optimal strategy through experience in the market; how does the restriction of information affect the ability to learn optimal behaviour; and what market structures lead to the best outcome in terms of efficient allocation of resources and maximizing the revenue for the seller?

These issues can be addressed theoretically through auction theory, experimentally through simulated auction environments with human volunteers competing in a controlled environment or by computer simulation with autonomous agents either following fixed strategies [9] or learning through interaction [6, 15, 13]. All approaches have a role to play in understanding auctions. Auction theory can tell us about the theoretical behaviour of agents in auctions and under certain certain constrained models can derive the optimal behaviour [11]. However, real world auctions rarely satisfy these constraints, and understanding of actual behaviour often requires observation and analysis of patterns of actual bidding. Experimental studies can help increase our understanding of how competitors behave in auctions through simulations of different auction structures with imaginary goods [8]. However, experimental studies suffer from drawbacks, chief among these being the difficulty in collecting an adequate amount of data and the problem with finding subjects with suitable experience in similar scenarios. Computer simulations offer an additional information source for addressing the issues that do not necessarily suffer the same problems.

This paper examines behaviour of autonomous adaptive agents (AAA) in simulations of common auction scenarios with a single seller. Recently, there has been much interest in studying the behaviour of agents in double auctions (auctions with multiple sellers and buyers) and in models of more constrained markets [4]. For example, Cliff [5] and Hu [10] have examined how alternative agent structures perform and the effect of market structure on performance. The aim of this research are to investigate how AAA behave in 'one-way' auction environments. The potential benefits of this project are: an economic test bed to examine effect on behaviour of alternative auction structures; an alternative mechanism to verify theoretical behaviour in auctions and compare with real world experimental studies; a new arena in which to study the evolution of multi-agent systems.

Although agent simulations offer a means of overcoming some of the problems with human experimentation, in order to be of use in the analysis of behaviour it is important to demonstrate that the agents can learn to perform well (in terms of meeting their objectives) and that they can learn the complexity of strategy exhibited by human competitors. Our approach to developing AAA for auctions is to start with a constrained environment where a single adaptive agent competes against agents that adopt the symmetric equilibrium strategy. This has the benefit that we can deduce the optimal strategy for the adaptive agent and thus assess performance and study the effects of alternative learning mechanisms on the quality of the strategy learnt.

Section 2 describes the auction environment. We adopt the private values model commonly used in auction theory and experimental studies [8]. The alternative agent learning mechanisms employed are described in Section 3. The learning mechanism used is similar to that employed by Cliff's ZIP agents [6], but with certain differences necessitated by the nature of the problem. The results for three alternative learning schemes are presented in Section 4. The first set of results in Section 4.1 demonstrate that the learning mechanism employed is able to learn the optimal strategy when provided with perfect information about what it should have bid after the event. This scenario is not realistic, since some information is usually withheld from the agent. However, it is desirable that any learning agent should be able to learn optimality in the complete information case, and the tests described in Section 4.1 give us a means of testing whether the agent meets the minimum level of performance we require. The second and third set of results contrast two approaches to learning when faced with uncertainty and find that the agent needs to utilise the information it is assumed to be party to under the PVM in order to reach the optimal strategy. Finally, the conclusions and future directions are discussed in Section 5.

We use the terms agents and bidders interchangeably. We use N to denote the number of agents including the adaptive agent and n to denote the number of non-adaptive agents (i.e. $n = N - 1$). We subscript the variables associated with the agents with i and those relating to a particular auction from a sequence of auctions with j, although we drop the second subscript unless explicitly required.

2 The Simulated Auction Model

2.1 Private Value Model

We use the private value model (PVM) proposed by Vicary [14] in the experimentation described in this paper. The PVM can be described in the following way. In an auction of N interested bidders, each bidder i has a valuation x_i of the single object. Each x_i is an observation of an independent, identically distributed random variable X_i with range $[0, \omega]$ (ω is the universal maximum price) and distribution function F.

The benefits of this model are that for certain auction mechanisms and assumptions (described in Section 2.2) there is provably optimal behaviour. Thus we have a clear way of gauging the performance of adaptive agents and assessing under what conditions learning is most efficient. This is a necessary condition to studying more interesting scenarios where the assumptions about the competitors behaviour do not hold true.

2.2 Auction Mechanisms

The four most used and studied auction forms are:

1. The open ascending price or *English* auction, where bidders submit increasing bids until no bidders wish to submit a higher bid;
2. The open descending price or *Dutch* auction, where the price moves down from a high starting point until a bidder bids, at which point the auction terminates.
3. The first-price sealed-bid auction (FPSB), where each bidder submits a single bid, the highest bidder gets the object and pays the amount he bid;
4. A second-price sealed-bid auction (SPSB), where each bidder submits a single bid, the highest bidder gets the object and pays the second highest bid.

Under the PVM, a Dutch auction is strategically equivalent to FPSB auction and an English auction is equivalent to a SPSB auction (strategic equivalence implies that for every strategy in one auction there is an equivalent strategy in the other auction that would have identical outcomes). Hence we restrict our attention to FPSB and SPSB auctions. The FPSB and SPSB simplify the bidding mechanism, since each agent can submit at most one bid for any auction. We denote the bids of the agents b_i. Each agent forms its bid with a bid function

$$\beta_i : [0, \omega] \rightarrow \Re_+, \qquad \beta_i(x_i) = b_i.$$

The set of all bids for a particular auction is denoted $B = \{b_1, b_2, \ldots, b_N\}$. For both FPSB and SPSB, the winning agent, w, is the highest bidder,

$$w = \arg\max_{i \in N}(b_i \in B).$$

The price paid by the winning agent, p, is dependent on auction structure. In a FPSB, the price paid is the highest bid,

$$p = \max_{i \in N}(b_i \in B).$$

In a SPSB, the price paid is the second highest bid,

$$p = \max_{i \in N, i \neq m}(b_i \in (B - b_m)).$$

where b_m is the largest bid, i.e. $b_m = \max_{i \in N}(b_i \in B)$.

2.3 Optimal Strategies

An agent's profit (or reward) is

$$r_i(x_i) = \begin{cases} x_i - p & \text{if } i = w \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

All agents are risk neutral, i.e. they are attempting to maximize their profit. The optimal strategy is the bid function that will maximize the profit of the agent. A symmetric equilibrium is a Nash equilibrium in which all bidders follow the same strategy. Hence a bid function, B^*, is a symmetric equilibrium if any one agent can do no better (in terms of maximizing expected reward) than follow B^* if all other agents use B^*.

First Price Sealed Bid The symmetric equilibrium strategy in a FPSB for agent i is the expected value of the largest of the other agents' private values, given that the largest of these values is less than the private value of agent i, i.e. the agent should bid as high as all the other agents' values (not bids) without exceeding its own value. More formally, let $Z_1, Z_2, \ldots Z_n$ be the order statistics of the agent values $X_1, X_2, \ldots X_n$. The symmetric equilibrium strategy for agent N is

$$\beta(x_N) = E(Z_n | Z_n < x_N) \tag{2}$$

The optimal strategy is dependent on the form and assumed commonality of the value distribution function F and the independence of the bidders' values. When F is a uniform distribution on $[0, 1]$ the symmetric equilibrium strategy is

$$\beta(x_i) = \tfrac{N-1}{N} x_i$$

Second Price Sealed Bid The symmetric equilibrium strategies in a SPSB auction are given by

$$\beta(x_N) = x_N. \tag{3}$$

The optimal strategy for a SPSB auction is independent of the form of the distribution function F and does not require that all bidders have the same value function. Proofs and a more complete description of auction formats are given in [11].

2.4 Auction Simulation Structure

The agents compete in a series of k auctions indexed by $j = 1, \ldots, k$. For any auction j, each bidder is assigned a value $x_{i,j}$ by sampling F. Prior to bidding, each bidder is aware of:

1. The realisation $x_{i,j}$ of X_i;
2. The distribution function common to all bidders, F;
3. The universal maximum value, ω;
4. The number of competitive bidders, N.

Once the auction is complete, the bidder is informed of whether they were the winner and the price the winner must pay. No other information relating to the other agents' bids is made available. As discussed in Section 3, this restriction of information affects an agent's ability to learn a good strategy. The agent is also unaware of the number of auctions in an experiment.

We use a uniform distribution on $[0, 1]$ for F and fix N for the duration of an experiment. All experiments involve a single adaptive agent competing against n non-adaptive agents following the optimal strategy relevant to the auction structure (Equations 2 and 3).

3 The Adaptive Agent Structure

With no loss of generality we assume the adaptive agent is bidder N. An adaptive agent attempts to maximize its expected reward, $E(r_N)$ where r_N is as defined in Equation 1, by finding an optimal strategy or bid function B^*. Under the model described in Section 2 the optimal strategy is as given in Equation 2 for FPSB and Equation 3 for SPSB auctions. The key issue is how the agent uses the information made available to it after the auction is complete in order to alter its behaviour to that closer to the optimal. Hence the agent needs to be able to assess or estimate how good a particular bid was and how that information should be employed to alter behaviour. We restrict the class of bid functions to linear functions of the form

$$B_N = x(1 - \mu_N),$$

where μ_N represents the percentage margin below its value that the agent bids at. This form allows us to simply relate the behaviour of the adaptive agent to the non-adaptive agents and measure how close the agent is to adopting the optimal strategy. The optimal strategies (and hence the strategies followed by the non-adaptive agents) are recreated by setting $\mu = 0$ for SPSB and $\mu = 1/N$ for FPSB. Thus the problem of learning a strategy is now the problem of learning the best margin μ based on the information available.

After an auction has been concluded the agent attempts to calculate the optimal bid it could have made, o_j, and hence the margin it should have used to bid in the auction, $d_{N,j}$, in order to maximize profit. It uses the estimate of the desired margin for auction j to adjust its margin using the Widrow-Hoff update rule

$$\mu_{N,j+1} = \mu_{N,j} + \sigma_{N,j}, \tag{4}$$

$$\sigma_{N,j} = \beta(d_{N,j} - \mu_{N,j}). \tag{5}$$

We drop the auction subscript j for the remainder of this Section for clarity. $b_1, ..., b_n$ are the bids of the n other agents and $y_1, ..., y_n$ are the bids sorted into ascending order, i.e. the observed values of the n bid order statistics $Y_1, ..., Y_n$. The optimum bid, o, is defined as the bid which would have maximised the agents reward r_N, given all the information about the other bids.

In auctions where the highest bid wins, the optimal bid is always equal to a small amount above the bid of the highest other bidder, as long as that bid is below the winning bid. If we ignore the small increment, then when $y_n < x_N$ the optimal bid is the largest of the other bids, $o = y_n$. Given the optimal bid, the optimal margin is

$$d_N = 1 - \frac{o}{x_N}. \tag{6}$$

There are two crucial issues to address in designing a learning mechanism for auction environments. The first issue is what the agent should do when the price paid is higher than the agent's value for that auction. In this scenario it is unclear what the optimal bid should have been, because the maximum attainable reward of zero is achieved by any bid less than the agent's value. Our initial approach is simply to set the optimal bid to the previous actual bid, thus not altering the bid margin. Although reasonable in the context of a single auction, this approach ignores the potential of using these bids to learn about the strategy of the other agents. However, as we demonstrate in Section 4.1, this approach is sufficient to learn optimal behaviour if the agent can calculate o accurately in all cases when $p \leq x_N$. The second issue is what to do when the optimal bid cannot be calculated exactly. Two approaches to this problem are presented in Section 3.1.

3.1 Estimating the Optimal Bid

Whether the agent can calculate the optimal bid exactly is dependant on the information available to the agent, which itself is dependant on the auction structure. In order to calculate o in a FPSB and SPSB auction, the agent needs to know the largest bid of the other agents. One of the key differences between FPSB and SPSB is the circumstances under which the agent does not know with certainty what the optimal bid should have been. The two cases where the agent faces this situation of incomplete information are

1. **Case 1**: in a FPSB when the agent wins ($p = b_N$, $w = N$) the agent is not told any of the other agents' bids. The adaptive agent only knows that the highest bid of the other agents, and hence the optimal bid, is in the interval $[0, p)$.
2. **Case 2**: In a SPSB when the agent loses ($p \leq b_N$, $w \neq N$), the agent is only informed of the second highest bid of the other agents, y_{n-1}, is less than or equal to p. The agent does know that the highest bid is on the interval $[p, \omega)$, and that its optimal bid is in the interval $[p, x_N]$.

In all other scenarios when $p < x_N$ the agent is able to determine exactly what the best bid would have been, and adjust its behaviour accordingly. In both cases of incomplete information, the agent is interested in estimating y_n, the highest bid (not counting the agents own) and hence the optimal bid. There are two approaches we investigate for estimating the optimal bid in the two cases of incomplete information.

1. *Naive estimation*: The agent simply guesses a value on the interval, thus in Case 1 the agent estimates y with a random value on $[0, p)$ for case 1 and $[p, \omega)$ in Case 2.
2. *Estimate based on assumptions*: Based on the PVM assumptions about the behaviour of the other agents, the agent can estimate the expected value of the relevant order statistics. This method is described in section 3.2

3.2 Estimating from Order Statistics

If we assume that each $b_1, ..., b_n$ is an independent indentically distributed ($i.i.d$) random variable, then we can define the problem facing the agent as calculating or estimating the expected value of the n^{th} order statistic conditional on the information provided by the auction. Under the PVM each bidders' value is an observation of $i.i.d.$ random variable, hence the assumption is equivalent to assuming that each agent is following an identical bidding strategy (i.e. the model is symmetric) and that this strategy remains constant from auction to auction (i.e. the agents are non-adaptive). These assumptions are common in auction theory.

Let f_b be the common density function of the bids and F_b the distribution function. The distribution of the n^{th} order statistic is

$$g_n(y_n) = n \cdot f_b(y_n) \cdot [F_b(y_n)]^{n-1}$$

Case 1: FPSB, $p = b_i$, $w = N$. In a FPSB auction where the agent wins (case 1), the optimal bid is the highest of the other bids, i.e. y_n. The distribution of the n^{th} order statistic conditional on $y_n < p$ is

$$g_n(y_n|y_n < p) = \frac{f_b(y_n)[F_b(y_n)]^{n-1}}{\int_{t=-\infty}^{p} f_b(t)[F_b(t)]^{n-1}dt}$$

and the agent wishes to calculate $E(y_n|y_n < p)$. See [7] for more background on conditional distributions of order statistics. Suppose we know f_b is uniform on $[0, \omega]$, then

$$g_n(y_n|y_n < p) = \frac{n[y_n]^{n-1}}{p^n}$$

and

$$E(y_n|y_n < p) = p\frac{n}{n+1}.$$

Hence, if the agent assumes the distribution of y_n is uniform, then the estimate of the optimal bid is

$$o = b_i\frac{n}{n+1}.$$

Case 2: SPSB, $w \neq N$. In a SPSB auction where the agent loses, we have two possible situations. Firstly when $p = b_N$, i.e. the adaptive agent's bid was the second largest. There are two pieces of information available, firstly that $y_n > p$ and secondly that $y_{n-1} < p$.Secondly when $p > b_N$, i.e. the adaptive agent's bid was not the second largest. The agent can now infer that $y_{n-1} = p$. Both scenarios require the distribution of the n^{th} order statistic conditional on the $Y_{(n-1)}$. From [7], the conditional distribution is

$$g_n(y_n|y_{n-1}) = \frac{f_b(y_n)}{[1 - F_b(y_{n-1})]}.$$

If we assume f_b is uniform on $[0, 1]$,

$$g_n(y_n|y_{n-1}) = \frac{1}{[1 - y_{n-1}]}$$

and

$$E(y_n|y_{n-1}) = \frac{1 + y_{n-1}}{2}.$$

If the agent assumes the distribution of y_n is uniform on $[0, 1]$ we can generalise the two scenarios so that the estimate of the optimal bid is

$$o = \frac{x_i + p}{2}.$$

4 Results

All experiments involve a single adaptive agent competing against 15 non-adaptive agents following the optimal strategy, over 10000 auctions. A learning rate of $\beta = 0.1$ was used.

4.1 Agents Learning with Sufficient Information to Estimate the Optimal Bid

Before continuing to examining performance in environments where the agent is confronted with *a posteriori* incomplete information, it is important to verify that the learning mechanism employed can at least converge to optimality when the agent is allowed to know the best it could have done. These "cheating" agents provide us with reassurance that the learning mechanism proposed in Equations 4, 5 and 6 will tend toward the optimal strategy with complete *a posteriori* information.

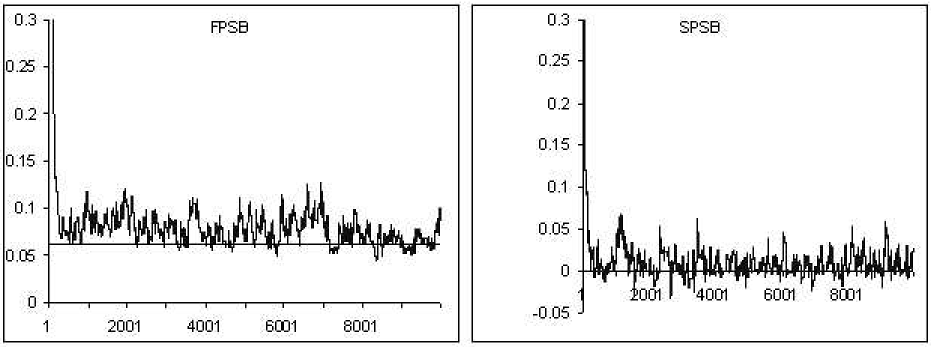

Fig. 1. "Cheating" adaptive agent bid margin for FPSB and SPSB auctions. The straight line for FPSB is the optimal bidding margin.

Figure 1 shows the bid margin $\mu_{N,j}$, of a "cheating" adaptive agent over a single experiment of 10000 auctions for FPSB and SPSB respectively. The optimal margin for FPSB is $1/N$ and for SPSB the optimal margin is 0. Table 1 shows the total profit made by the adaptive agent and the 15 non-adaptive agents over the last 5000 auctions in a run of 10000.

Table 1. Total profit over the last 5000 auctions for the cheating agent against 15 non-adaptive optimal agents in FPSB and SPSB auctions

FPSB					SPSB				
Adaptive Agent	Non-Adaptive Agents				Adaptive Agent	Non-Adaptive Agents			
Profit	Min	Max	Mean	Std Dev	Profit	Min	Max	Mean	Std Dev
17.93	16.70	20.73	18.59	0.98	19.44	16.59	21.81	18.51	1.39

Figure 1 indicates that there may be some minor deviation from optimality with FPSB auctions. However, this small deviation is not significant enough to effect the profits of the adaptive agent which, as can be seen in Table 1. For both FPSB and SPSB the adaptive agent profit is within the range of the profit of the other agents, and is not significantly less than the mean. The cheating agent is quickly learning the optimal strategy for SPSB.

Hence the learning mechanism describe in Section 3 converges to the symmetric equilibrium strategy when provided with complete information. The next two sets of experiments are designed to see how the adaptive agent performs in the more realistic scenario of incomplete information, which requires that the optimal bid be estimated rather than known.

4.2 Agents using Naive Estimation

Figure 1 and Table 1 show that the agent can learn an optimal strategy for FPSB and SPSB when given the optimal bid in the case when $p \leq x_N$. However, in both auction types this information is not always available. Section 3.1 describes two approaches adopted faced with this uncertainty. The most obvious strategy when faced with uncertainty as to the optimal bid, the naive approach, is simply to guess a value on the interval in which the optimal bid must lie. Figure 2 shows that adopting this approach does not lead to the optimal strategy. This point is reinforced by Table 2 which clearly show the adaptive agent is making significantly less profit than the adaptive agents and is doing markedly worse in FPSB than SPSB.

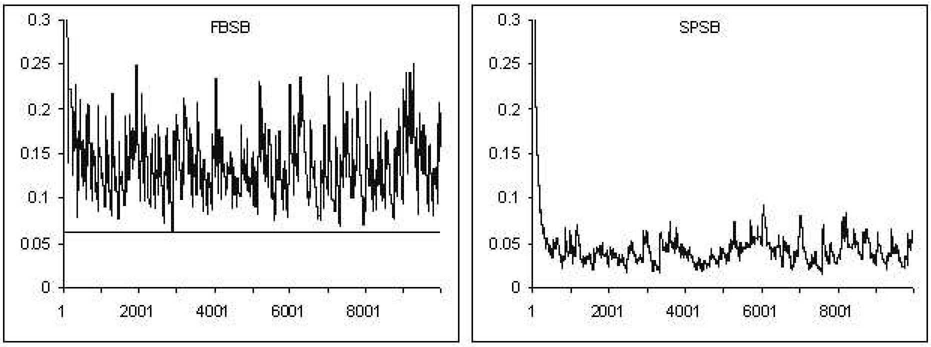

Fig. 2. Naive adaptive agent bid margin FPSB and SPSB auctions.

Table 2. Total profit over the last 5000 auctions for the naive agent against 15 non-adaptive optimal agents in FPSB and SPSB auctions

FPSB					SPSB				
Adaptive Agent	Non-Adaptive Agents				Adaptive Agent	Non-Adaptive Agents			
Profit	Min	Max	Mean	Std Dev	Profit	Min	Max	Mean	Std Dev
10.77	17.20	20.95	19.20	1.19	16.74	16.96	20.62	19.06	1.11

Failure to reach the optimum strategy can be explained by the uniform sampling method of estimating the optimal bid. In a FPSB, when the agent wins the estimate is made on the range $[0, p]$, and will on average be $\frac{p}{2}$. However, the expected value of the second highest bid is $\frac{N-1}{N}p$, hence the agent is repeatedly increasing it's bid margin by too much, hence losing auctions it could have won with a profit. In SPSB the estimate is uniform on the range $[p, x_N]$. The improved performance in SPSB shown by the greater profit is due to the fact that the range $[p, x_N]$ will be smaller than the range $[0, p]$, hence the bias introduced by the random estimation will be less.

4.3 Agents using Estimation based on Assumptions

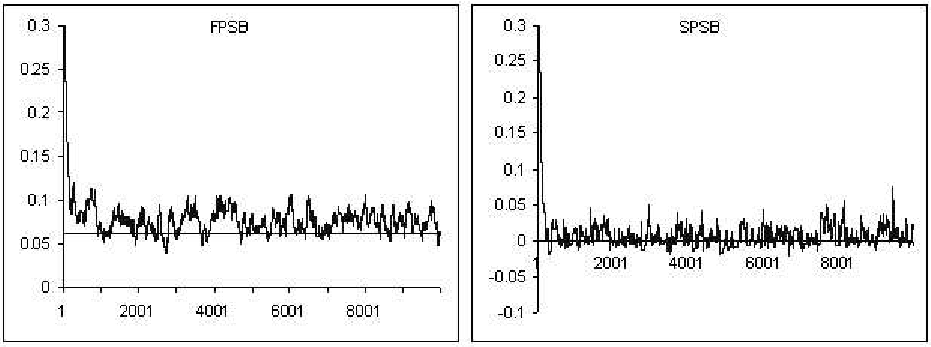

Fig. 3. Estimating adaptive agent bid margin for FPSB and SPSB auctions.

Table 3. Total profit over the last 5000 auctions for the estimating agent against 15 non-adaptive optimal agents in FPSB and SPSB auctions

FPSB					SPSB				
Adaptive Agent	Non-Adaptive Agents				Adaptive Agent	Non-Adaptive Agents			
Profit	Min	Max	Mean	Std Dev	Profit	Min	Max	Mean	Std Dev
17.86	16.51	20.66	18.62	1.12	19.13	15.81	22.81	19.02	2.06

Figure 3 and Table 3 show the results for the estimating agent. The results indicate that the adaptive agent is doing as well as the non-adaptive agents in terms of profit. The agent margin converges towards the optimal, with wider variation for FPSB than SPSB. This demonstrates that the learning mechanism is able to utilise the information assumed available under the PVM to learn an optimal strategy.

5 Conclusions and Future Directions

This paper has introduced an agent model of common auction scenarios, FPSB and SPSB, using the private values model and examined how an autonomous adaptive agent using a basic learning mechanism can learn to behave optimally in terms of evolving towards a symmetric equilibrium solution. We think it it is a necessary prerequisite that any learning mechanism should be at least able to learn an optimal strategy given the necessary information to do so. We have demonstrated that the learning mechanism described in Section 3 is sufficient for the agent to make as much profit as the non-adaptive agents in both FPSB and SPSB auctions. In order to cope with the uncertainty inherent in accurate auction simulations, the agent using this learning mechanism must estimate in certain situations dependent on auction structure. We have presented two techniques for this estimation, a naive approach which was found to be insufficient and an estimation technique built on the assumptions of the PVM which quickly converged to optimality. This work will be extended through the following stages.

5.1 Future Work

We will continue to experiment with the PVM in order to find the most robust learning mechanism for scenarios with a known optimal strategy. This process will include: a thorough analysis will be conducted of the apparent slight deviation from optimality suggested by Figure 1; an extension of the learning mechanism that uses the information available in situations where the price paid is above the agents value; an assessment of how the learning mechanisms perform with alternative value distributions; an assessment of how the learning mechanisms perform when the number of adaptive agents can vary during an experiment; relaxation of the PVM assumptions. Having developed learning mechanisms best able to cope in situations where optimality can be measured, we will then examine how these mechanisms perform when the assumptions are no longer met and begin an investigation into behaviour of multi-adaptive agent models of auction scenarios.

References

1. FreeMarkets Inc. `http://www.freemarkets.com/`.
2. FreightTraders Ltd. `http://www.freight-traders.com/`.
3. A.Chavez, A. Moukas, and P. Maes. Challenger: A multi-agent system for distributed resource allocation. In W. Lewis Johnson and Barbara Hayes-Roth, editors, *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, pages 323–331, New York, 1997. ACM Press.
4. A. J. Bagnall and G. D. Smith. An adaptive agent model for generator company bidding in the uk power pool. In *Lecture Notes in Computer Science 1829, Artificial Evolution*, pages 191–203, 1999.
5. D. Cliff. Evolution of market mechanism through a continuous space of auction-types. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC2002)*, pages 206–209, 2002.
6. D. Cliff and J. Bruten. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets, 1997.
7. H. A. David. *Order Statistics*. John Wiley and Sons, 1970.
8. D. Friedman and J. Rust, editors. *The Double Auction Market*. Perseus Publishing, 1991.
9. D. Gode and S. Sunder. Allocative efficiency of markets with zero intelligence traders. *Journal of Political Economy*, 101:119–137, February 1993.
10. Junling Hu and Michael P. Wellman. Multiagent reinforcement learning: theoretical framework and an algorithm. In *Proc. 15th International Conf. on Machine Learning*, pages 242–250. Morgan Kaufmann, San Francisco, CA, 1998.
11. V. Krishna. *Auction Theory*. Harcourt Publishers Ltd, 2002.
12. M. Sawhney and S. Kaplan. The-b-to-b-boom, let's get vertical. *Business 2.0*, September 1999.
13. G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the Third acm Conference on Electronic Commerce*, pages 206–209, 2001.
14. W. Vickery. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
15. Michael P. Wellman and Junling Hu. Conjectural equilibrium in multiagent learning. *Machine Learning*, 33(2-3):179–200, 1998.