

A Heuristic Bidding Strategy for Multiple Heterogeneous Auctions

Patricia Anthony
Dept. of Electronics and Computer Science
University of Southampton
Highfield, Southampton, SO17 1BJ, UK.
pa99r@ecs.soton.ac.uk

Nicholas R. Jennings
Dept. of Electronics and Computer Science
University of Southampton
Highfield, Southampton, SO17 1BJ, UK.
nrj@ecs.soton.ac.uk

ABSTRACT

Online auctions are increasingly being used as a medium to procure goods and services. As the number of auction sites increases, however, consumers will inevitably want to track and bid in multiple auctions (with multiple protocols) in order to get the best deal for their desired goods. To this end, this paper reports on the development of a heuristic decision making framework that an autonomous agent can exploit to tackle the problem of bidding across multiple heterogeneous auctions. The framework enables the agent to adopt varying tactics and strategies that attempt to ensure that the user's objectives are satisfied. Through empirical evaluation, the agent's performance is shown to be effective even when there are multiple such agents in the environment at the same time and when the agent cannot accurately determine the type of environment that it is situated in.

Keywords

multiple auctions, bidding strategy, intelligent agents.

1. INTRODUCTION

Online auctions are one of the most popular and effective ways of trading goods over the Internet; indeed there are more than 2500 sites that operate online auctions¹. These auction houses conduct many different types of auctions according to a variety of protocols including English, Dutch, first-price sealed bid and second-price sealed bid (Vickrey). Moreover, the auctions operate with varying start and end times and many of them are likely to be active at any one moment in time. For these reasons, consumers are faced with a difficult task when it comes to selecting and bidding in these auctions. In order to purchase a particular item, they need to monitor multiple auctions at multiple sites before selecting the one to participate in. Even after making this selection, they still have to decide on the amount to bid in order to get the desired item under conditions that are consistent with their preferences.

To assist users, many sites now offer simple bidding proxies that will bid in a stated auction up to some predefined limit (e.g. eBay and Yahoo!Auction). However, such systems typically have the problem of only being able to operate

at a single auction site or only with a particular auction protocol. Users can also utilise the services of auction search engines (such as BidXS² and AuctionBeagle³). However, while these facilities allow users to monitor the progress of multiple concurrent auctions, all the bidding decisions are left to the human users.

To address these shortcomings, we believe it is necessary to develop an autonomous agent that can participate in multiple heterogeneous auctions, that is empowered with trading capabilities, and that can make purchases autonomously [3, 1]. In more detail, the agent should monitor and collect information from the ongoing auctions, make decisions on behalf of the consumer and endeavour to guarantee the delivery of the item. The agent must ensure that it never bids above the private valuation (the maximum amount that the consumer is willing to pay) and it should try to get the item in a manner that is consistent with the consumer's preferences (e.g. at the earliest time, at the lowest price, or with maximum chance of succeeding).

To this end, this paper reports on our work in developing such a bidding agent. The agent has a range of strategies that it can employ depending on the user's aims and the environment in which the agent finds itself. Here, we consider three different types of protocol: English, Dutch and Vickrey⁴. The strategies themselves are heuristic in nature because the multiple heterogeneous auctions environment is highly complex, dynamic and unpredictable, making it impossible to find an optimal strategy that can be used in practical contexts [5]. Moreover, our early investigations showed that the effectiveness of the strategies was heavily influenced by the nature of the environment [3]. For this reason, we decided to have different strategies for different circumstances. As the range of potential strategies is huge, we decided to use a genetic algorithm (GA) to search for effective strategies for each of the various environments that we identified [2]. We chose this particular method because GAs have been shown to perform well in areas where the search space is large and not well understood [9]. Having evolved the strategies (offline), the agent adopts the one that is most appropriate to its prevailing context (online).

¹<http://www.internetauctionlist.com>

²<http://www.bidxs.com>

³<http://www.auctionbeagle.com>

⁴The first-price sealed bid is not considered here because of its similarities to the Dutch auction [8].

Against this background, this paper advances to the state of the art in the following ways. Firstly, we extend the work presented in [2] to evolve the complete set of strategies for our trading scenarios (in [2] we merely described the method and showed that it was successful in a very limited range of environments). These strategies consist of multiple evolved sub-behaviours that are appropriate in different environment settings with different objectives. This collection of strategies for a single agent is termed *the intelligent bidding strategy* in the remainder of this paper. Secondly, we evaluate the intelligent bidding strategy and show that it is effective across a broad range of scenarios (including those where the agent cannot accurately assess the type of environment it is situated in). Thirdly, given the success of our bidding strategy, we believe that it may be widely adopted. Thus, we present results about what happened to the performance of the buyers and sellers and the market efficiency when multiple agents adopt this strategy in a given marketplace.

The remainder of this paper is structured in the following manner. In the next section, we briefly outline the design of our intelligent bidding strategy. Section 3 reports on the evaluation of the strategy in various settings. Section 4 discusses related work and finally Section 5 concludes.

2. DESIGNING THE BIDDING STRATEGY

Before describing the decision-making framework, it is necessary to detail our assumptions about the environment. Firstly, we consider three auction protocols: English, Dutch and Vickrey (three of the most common types). Secondly, all auctions have a known start time and English and Vickrey auctions have a known end time. Thirdly, our bidding agent is given a hard deadline (t_{max}) by when it must obtain the desired item. It is also told the consumer's private valuation (p_r) for this item and the consumer's intention (either looking for a bargain, desperate for the item or a combination of the two). Finally, the agent must not buy more than one instance of the desired item. We use a simulated electronic marketplace to evaluate our bidding strategies (see [3, 1] for more details). This marketplace consists of a number of English, Dutch and Vickrey auctions that run concurrently. There are several bidders participating in each auction. They operate in a single auction, have the intention of buying the target item and possessing certain behaviour based on the type of auction that they are participating in. These bidders' strategies are based on the dominant strategies of the respective one shot single auctions [8].

2.1 The Bidding Strategy Framework

The agent's decision making model works in the following manner (as summarised in Figure 1). The bidder agent builds an active auction list (auctions that have started but not reached their end times, denoted $L(t)$) and gathers relevant information (e.g. start and end times, current bid values) about them. It then calculates the *current maximum bid* it is willing to make at the *current time*. The current maximum bid, by definition, will always be less than or equal to the private valuation. To determine the current maximum bid, the agent considers several *bidding constraints* including the remaining time left, the remaining auctions left, the user's desire for a bargain and the user's level of desperation. For each such bidding constraint,

```

while ( $t \leq t_{max}$ ) and (Item_Is_Not_Acquired)
  Build active auction list;
  Calculate current maximum bid using the agent's strategy;
  Select potential auctions to bid in, from the active auction list;
  Select target auction as one that maximises agent's expected utility;
  Bid in the target auction using current maximum bid as reservation
  price at this time;
endwhile

```

Figure 1: Bidding Agent's Top-Level Algorithm

there is a corresponding function that suggests the value to bid based on the constraint at that time. These functions (based on [6]) are parameterized by two key values: k (range [0..1]) is a constant that determines the value of the starting bid and β (range [0.005..1000]) defines the shape of the curve (and so the rate of concession to p_r). As an example, when $k \geq 0.5$ and $\beta \geq 1$, the agent demonstrates a reasonable degree of desperation and starts bidding close to p_r and quickly reaches p_r . At the other extreme, the agent can demonstrate hard bargaining behaviour ($k < 0.5$ and $\beta < 1$), where it makes a low initial bid and only concedes in a very slow fashion. All behaviours in between are also possible by setting the parameters appropriately. At any given time, the agent may consider any of the bidding constraints individually or it may combine them depending on the situation (what the agent sees as being important at that time). If the agent combines multiple bidding constraints, it allocates a weight to each of them to denote their relative importance. The set of functions is referred to as the *tactics* and the combination of these tactics are referred to as the *strategy*.

Figures 2 and 3 illustrate how these bidding constraints can be combined to produce a single bid value for a given time. In particular, Figure 2 shows the behaviour of each bidding constraint. This behaviour suggests how the agent should behave when only a single constraint is considered. For example, if the agent considers desire for a bargain as the only bidding constraint, the it should start bidding at a low bid value and should slowly move toward its private valuation. On the other hand, if the agent wishes to consider the remaining time left as the single criterion, then it should start bidding at a reasonably high value and slowly reach its private valuation. Figure 3 show various combinations of the bidding constraints using weighted combinations. In this case, (0.25, 0.25, 0.25, 0.25) indicates that equal weighting has been applied to all the bidding constraints, (0.50, 0.00, 0.50, 0.00) indicates that the agent only considers bidding constraints of the remaining time left and the desire for a bargain and places equal weight of 0.50 for both bidding constraints. It can also be seen that by varying the weights of the bidding constraints, different bidding patterns can be generated as shown in Figure 3. Based on the value of the current maximum bid, the agent selects the potential auctions in which it can bid and calculates what it should bid at this time in each such auction. The auction and the corresponding bid with the highest expected utility is then selected from the potential auctions as the target auction.

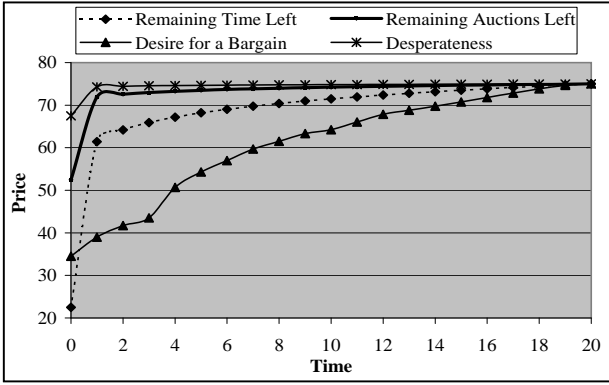


Figure 2: The Bidding Constraints' Behaviour

Finally the agent bids in the target auction. This process is repeated until the item is acquired or until the given time is reached.

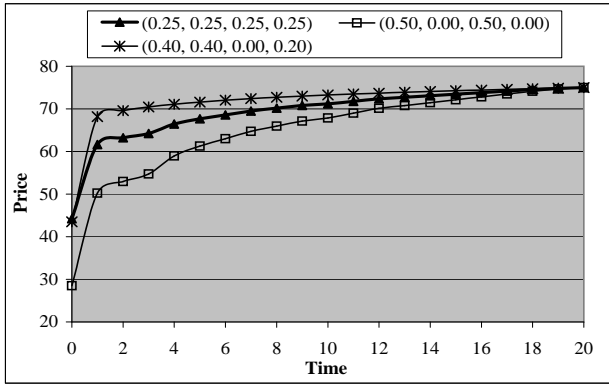


Figure 3: Various Combinations of the Bidding Constraints

Our early experiments in this domain led to several conclusions [3]. Firstly, p_r is one of the most important factors that needs to be considered when determining the strategy that should be employed by the agent. For example if p_r is low, the probability of winning the auction is low and the agent is unlikely to be able to get a bargain, whereas if p_r is high, the agent has a better chance of obtaining the item with a bargain. Secondly, the strategies to be used by the agent need to be tailored to the prevailing context, since not all strategies work well in all situations. Thus, a successful strategy in one situation may perform badly in another. Nevertheless, it was possible to determine that certain classes of strategy are effective in environments that have particular characteristics. In this case, the key defining characteristics of an environment were found to be the number of auctions that are active before t_{max} . Given this, it was decided to evolve strategies that are effective for environments classified in this way.

2.2 Evolving Strategies

The performance of the bidding agent is heavily influenced by the strategy employed, which, in turn, relates to the values of k and β in a given tactic and the weights for each tactic when they are combined. The number of strategies that can

RP1	FE1	STLA	RP2	FE1	STLA	RP3	FE1	STLA
		STMA			STMA			STMA
		MTLA			MTLA			MTLA
		MTMA			MTMA			MTMA
		LTLA			LTLA			LTLA
	LTMA	LTMA		LTMA				
	FE2	STLA		FE2	STLA		FE2	STLA
		STMA			STMA			STMA
		MTLA			MTLA			MTLA
		MTMA			MTMA			MTMA
		LTLA			LTLA			LTLA
	FE3	STLA		FE3	STLA		FE3	STLA
		STMA			STMA			STMA
		MTLA			MTLA			MTLA
		MTMA			MTMA			MTMA
LTLA		LTLA	LTLA					
LTMA	LTMA	LTMA						

RP1: Low Private Valuation

RP2: Medium Private Valuation

RP3: High Private Valuation

FE1: Fitness Equation I

FE2: Fitness Equation II

FE3: Fitness Equation III

STLA: Short Time Less Auctions

STMA: Short Time Many Auctions

MTLA: Medium Time Less Auctions

MTMA: Medium Time Many Auctions

LTLA: Long Time Less Auctions

LTMA: Long Time Many Auctions

Table 1: The Environments

be employed is infinite, so, therefore, is the search space. Here, we decided to use GAs to search offline for the most successful strategies in the predefined environments (see [2] for more details). The individuals in the populations are the bidding agents and their genes consist of the parameters for the four different tactics and the relative weight for each of them. To compute the fitness function, we consider three plausible alternatives. These are the individual success rate in obtaining the item (FE1) and two variations based on the average utility. In the first case (FE2), the agent gets a utility of 0 if it fails to obtain the item. The final utility function (FE3) is similar to FE2 but the individual is penalised if it fails to get the item. FE1 is used if delivery of the item is of utmost importance, FE2 if the agent is looking for a bargain, and FE3 if delivery of the item and looking for a bargain are equally important.

In [2] we showed, through empirical evaluation, that successful strategies with the obvious behaviours (as dictated by the various fitness equations), could be evolved for a very limited set of environments. The next step, and the main contribution of this paper, is to complete the evolutionary process for all of the target environments and to show that the ensuing agent, composed of the evolved sub-behaviours, would actually perform effectively across a wide range of environments.

2.3 The Intelligent Bidding Strategy

Having shown that GAs can effectively evolve strategies for different environments, the final step is to combine this knowledge into a single intelligent bidding agent. The agent has at its disposal knowledge about which strategies are effective in which environments and assuming it can assess the environment accurately, it simply has to deploy the appropriate strategy.

In more detail, we subdivided our trading environments into 54 categories (shown in Table 1) based on the key determinants that influenced the agent’s bidding decision (namely, the private valuation, the remaining time left, the remaining auctions left and the agent’s behaviour (whether looking for a bargain, desperate or a mixture of both)). This categorisation of the classes of environment is important because it enables the agent to tune its bidding strategy to its prevailing circumstances. The agent’s private valuation can be broadly categorised by its value: low (RP1), medium (RP2) or high (RP3). These three categories are then refined further into three sub-categories based on the agent’s behaviour (FE1 is desperate, FE2 is looking for a bargain and FE3 aims for a balance of both). Finally, each sub-categorisation is further divided into groupings that consist of environments in which there are short time less auctions (STLA), short time many auctions (STMA), medium time less auctions (MTLA), medium time many auctions (MTMA), long time less auctions (LTLA) and long time many auctions (LTMA) (see [3] for more details). As an example, RP1FE2STMA represents a strategy that was evolved in an environment with a low private valuation, where the bidding time is short, where there are many auctions and where the consumer is interested in a bargain.

The categorisation of the private valuation is made based on the auction closing price distribution (here, the closing price mean is 76 and the standard deviation is 5). Thus, 50% of the auctions should be won by bidders with medium private valuations, 25% by bidders with low private valuations and the remaining 25% by bidders with high private valuations. In real market settings, the price of each desired item will naturally vary depending on the type of the item itself (e.g. a diamond ring usually costs more than a book). However, the value of a given item can be directly mapped to the reservation price categorisation merely by obtaining the mean price of the item. This can be achieved using comparison price data from sites such as PriceSCAN⁵, DealTime⁶, and BottomDollar⁷. From this, the agent can calculate the mean price of a given item and regenerate price ranges for the low, medium and high private valuations. The subdivisions of the short time, medium time, long time, less auctions and many auctions can also be carried out in a similar manner.

We then used the search algorithm defined in Section 2.2 to evolve the best strategy for each environment. Thus, the agent gets the user’s private valuation, the item to be purchased, when it is required and the intention of the user (either looking for a bargain, desperate or some combination of the two). With this knowledge, the agent enters the marketplace and determines the number of active auctions in which it can participate within the given time constraint. Based on this combination of information, the agent determines which strategy to use in each auction round. This decision is captured in a rule base that maps the prevailing context to the strategy that has been evolved for that situation. Upon selection of the appropriate strategy the agent proceeds as defined in Figure 1.

⁵<http://www.pricescan.com/>

⁶<http://www.dealtime.com>

⁷<http://www.bottomdollar.com/>

3. EXPERIMENTAL EVALUATION

Here we are interested in two main types of experimentation. Firstly to show that our intelligent bidding strategy performs effectively in a wide range of bidding contexts. Secondly, given its success, we believe that it will be widely adopted and thus we want to understand what will happen when there are multiple such agents in the environment.

In the first set of experiments, the performance of the agent is measured in terms of success rate and average utility. The success rate is defined as the number of times, as a percentage, the agent is successful in obtaining the item. The utility of winning in an auction i is computed as $U_i(v) = (\frac{p_r - v}{p_r}) + c$, where v is the winning bid and c is an arbitrary constant 0.001 to ensure that the agent receives some value when the winning bid is equivalent to its private valuation.

In order to assess the robustness of our bidding agent, we wanted to consider how it would perform if it could not accurately determine the type of environment it is in (e.g. if the actual environment is RP1FE2STMA, what would happen if the agent believed it was RP1FE2MTMA and played the corresponding evolved strategy). This agent we label as *inaccurate* (c.f. the accurate agent that correctly determines its environment). In this particular experiment, the inaccurate agent always picks a strategy that is very close to the actual environment based on the bidding time⁸. As an example, if the actual environment is RP1FE2STMA, the inaccurate agent may believe that it is in RP1FE2MTMA and will use the evolved strategy for RP1FE2MTMA as its strategy. The mapping from the actual environments to the inaccurately classified environments is shown in Table 2. Here, if the actual environment is categorised as a short time (e.g. RP1FE1STMA), then the inaccurate environment is mapped to a medium time environment (e.g. RP1FE1MTMA).

As a control model, we use agent C that has a single fixed strategy based on the user’s behaviour. In particular, we picked the strategy that was evolved for the environments RP2FE1MTMA, RP2FE2MTMA and RP2FE3MTMA⁹. Thus, C has three different strategies; when a user is interested in a bargain, C employs a strategy that was evolved for RP2FE2MTMA, when a user is desperate for the item, C uses a strategy that was evolved for RP2FE1MTMA and when a user is looking for a combination of both, it utilises the strategy evolved based on RP2FE3MTMA.

The agents and the control model were run 1000 times in the marketplace (independently)¹⁰. Analysis of variance (ANOVA) was used to test the hypothesis about the differences of the success rate means (when running the experiment 200,

⁸It is also possible to select a strategy that is close to its private valuation or the number of active auctions or some combinations of the three.

⁹We chose these environments because the average time to procure the good in the marketplace is 50, the average total number of auctions ($|L(t)|$) is 30 and the average closing price is 76 and these values fall into the RP2MTMA category.

¹⁰This means that the intelligent agent and the control models are not competing against each other. Rather, they are competing against the simulated auction participants as described in Section 2.

Actual Environment	Inaccurate Prediction
Short	Medium
Medium	Short
Long	Medium

Table 2: The Actual Environments’s Mapping to the Inaccurate Environment

400, 600, 800 and 1000 times) and the procedure revealed that for all experiments, the differences between means were not significant ($F_{4,15} = 0.134, p > 0.05$) and thus the results obtained are statistically significant. The user’s requirement was randomly allocated where the user’s private valuation ranges from 70 to 82 and the time allocated ranges from 10 to 100. The number of auctions running in the marketplace is between 3 and 60 and there are between 2 and 10 participants in each such auction.

The performance of the intelligent agents (accurate and inaccurate) and the control model in terms of success rate is shown in Figure 4. The experiment is divided into four groups. The first three groups show the detailed performance of the agents and the control model based on a single behaviour (desperateness, bargain, and both) and the last group shows the overall performance when all three behaviours are considered¹¹. It can be seen that the accurate agent achieved the highest success rate in all groups. This shows that by having the ability to change the strategy according to the user’s preference and the environment it is situated in, our agent can maximise its chances of succeeding. This is different from C , in which a fixed strategy is used and where the agent views the environment and the user’s preference as static. In this case, C was only successful in the environment for which it was evolved. Even though the inaccurate agent’s success rate is lower than the success rate achieved by the accurate agent, it is still better than that of C . This is because the strategy selected by the inaccurate agent is one that is usually fairly close to the actual environment and this proximity means the deployed strategy is likely to be reasonably effective.

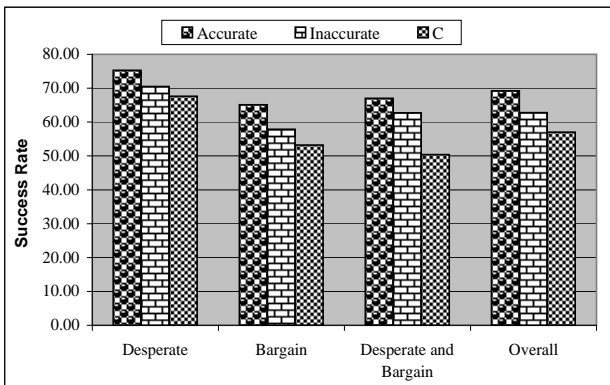


Figure 4: Success Rate Comparisons

¹¹Here, the user intention is generated randomly.

Figure 5 shows the average utility obtained by the intelligent bidding agents and the control model. Again, it can be seen that the accurate agent performed well compared to the other two models (except in the case of desperate behaviour where it obtained a lower average utility than the inaccurate one). This can be explained by the technique used by the inaccurate agent when selecting its strategy. As an example, when the actual environment is RP2FE1STMA, the inaccurate agent uses RP2FE1MTMA which is characterised by a low initial bid value and slowly reaching its p_r . The strategy for RP2FE1STMA is similar, but the rate at which it reaches p_r is faster. As a result, when the inaccurate agent obtains the item, the price it pays is low and this leads to a higher utility. Again, the inaccurate agent recorded a higher average utility than C but lower than the accurate agent. This indicates that, even when our agent cannot accurately assess the environment that it is in, by approximating its environment its performance is still better than an agent that uses a fixed strategy.

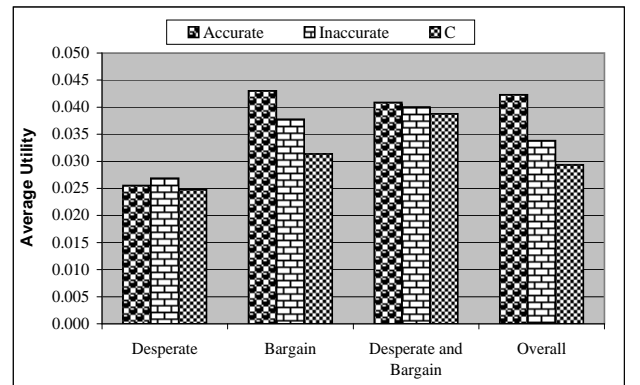


Figure 5: Average Utility Comparisons

Based on the level of achievement attained by the inaccurate agent, we conducted another experiment that analyses the performance of our intelligent bidding agent when it makes varyingly accurate predictions of the environment that it is in. This is achieved by gradually increasing the percentage of errors that it made when assessing its environment. Here, the set up of the experiment is as before but this time we focused on the overall behaviour of the agent. From Figure 6, it can be seen that as the error’s percentage increases, the success rate and the average utility obtained by the intelligent agents decreases. With a 50% error rate the agent’s success rate and average utility performances decrease by 12% and 13% and with a 100% error, its performance deteriorated by more than 25%. This result shows that even when our agent makes the wrong assessment half of the time, it is still able to achieve a reasonable success rate and a reasonable average utility (it obtained 61% success rate and 0.03291 average utility, whereas C obtained 57% success rate and 0.2934 average utility in the previous experiment).

We now turn to the second set of experiments. The purpose of these is to observe the effect on buyers and sellers (in terms of market efficiency) of having multiple intelligent bidding agents in the system. Here, we take a fixed environment consisting of 30 auctions (10 of each type) with 10 bid-

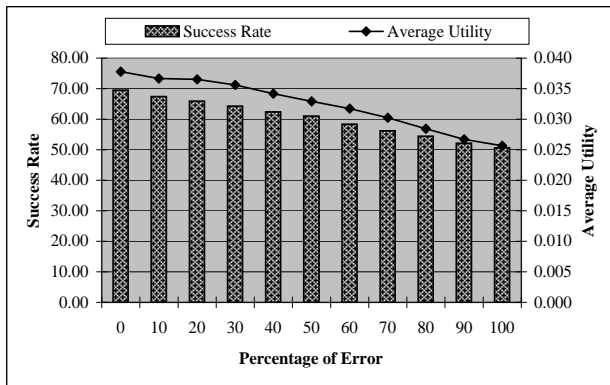


Figure 6: Agent's Performance with Varying Environment Prediction Accuracy

ders in each auction. The first group (called the *dummy bidders*) are those that operate in a single auction and have the intention of buying the target item and possessing certain behaviour (bargain seeking, desperate, etc.). They maintain information about the target item they wish to purchase, their private valuation of the item, the starting bid value and their bid increment. These values are generated randomly from a standard probability distribution. Their bidding behaviour is determined based on the type of the auction that they are participating in and are based on the dominant strategies of the respective one-shot single auctions (as described in Section 2). The second group is made up of the intelligent agents. These agents have all their individual and environmental parameter set in the same way as the dummy agents. Specifically, Figure 7 shows the supply and demand curves for the market. The intersection of the two curves is called the *equilibrium price* (when the supply of goods meets the demand), P_0 , and the quantity traded at this price is the *equilibrium quantity*. In this particular experiment, the equilibrium price is at 81 and the equilibrium quantity is 25.

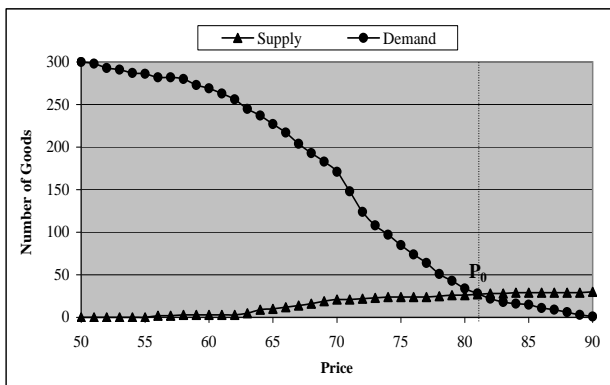


Figure 7: Supply and Demand

To measure the effectiveness of the market, we calculate its allocative efficiency and the Smith's Alpha (α) coefficient [11]. Allocative efficiency is defined as the total actual profit earned by all the sellers divided by the maximum total profit that could have been earned in an ideal market (expressed as a percentage). The Smith's Alpha coefficient measures

Number of Intelligent Agents	Intelligent Agents		Dummy Bidders	
	Average Utility	Average Number of Auctions Won	Average Utility	Average Number of Auctions Won
0			0.003419	26.85
30	0.007401	5.00	0.002807	20.70
60	0.006836	9.90	0.002593	16.60
90	0.006180	12.20	0.002509	14.65
120	0.004716	13.00	0.002490	11.80
150	0.004584	15.65	0.002254	9.45
180	0.004507	18.50	0.002090	6.50
210	0.004343	19.30	0.001846	4.70
240	0.003819	20.10	0.001410	2.55

Table 3: The Bidder's Average Utility

how close the actual trade prices are to the equilibrium, $\alpha = 100 * \sigma_0 / P_0$, where σ_0 is the standard deviation of trade prices around P_0 . We initially ran the marketplace with no intelligent agents and observed the average utility obtained by the dummy bidders. We then replaced 30 dummy bidders with 30 intelligent agents and observed the change in average utilities for both the dummy bidders and the intelligent agents. To be more precise, at each stage of the experiment, one dummy bidder in each auction is removed from the marketplace and is replaced with a single intelligent agent. This number is gradually increased in increments of 30 until there are 60 dummy agents and 240 intelligent agents¹².

Given this setup, Table 3 shows the average number of auctions won¹³ and the corresponding average utilities gained by the dummy bidders and the intelligent bidders for a total of 20 runs. It can be seen that the dummy bidders achieved the highest average utility when there were no intelligent agents (as expected). This value steadily decreases as more intelligent agents enter the system. The intelligent agents achieved the highest average utility when there were 30 such agents participating in the auctions. This is because there is only a small number of intelligent agents participating in the marketplace. This situation allows them to be in direct competition with the dummy bidders. After this point, the average utility gradually decreases as more intelligent agents enter the marketplace. These agents also obtain higher success rates as their number increases (as reflected in the number of auctions won by them). This is to be expected since there are now more intelligent agents dominating the marketplace leading to a scenario where they are now competing against each other (rather than competing against the dummy bidders) to maximise their user's satisfaction. It is interesting to note that even when there were 240 intelligent agents in the marketplace, the average utility obtained by these agents surpassed the highest average utility recorded by the dummy bidders. This implies that even when there

¹²Our auctions require a minimum of 2 agents to be present before they start and so we cannot move to only intelligent agents without distorting the experimental setting of running a constant numbers of agents.

¹³Not all auctions result in a sale (because of a lack of appropriate bids) and so the number of auctions won does not always sum to 30.

are many intelligent agents in the marketplace, they will still be able to achieve higher average utilities than those that utilised dominant strategies based on one shot auctions.

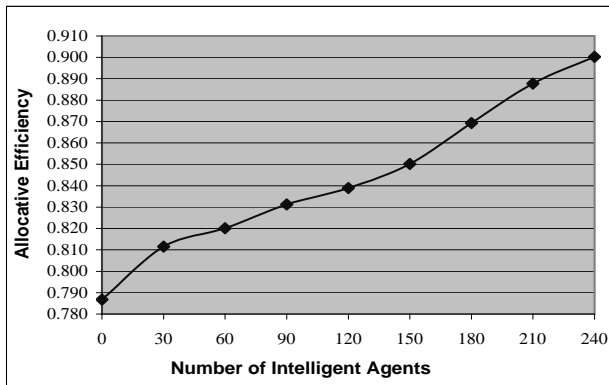


Figure 8: Relationship between Number of Intelligent Agents and Allocative Efficiency

The allocative efficiency and Smith's Alpha coefficient in relation to the number of intelligent agents in the market are shown in Figures 8 and 9. It can be seen that the allocative efficiency increases as the number of intelligent agents in the marketplace increases. We can also observe that the allocative efficiency rises from 0.79 (when there are no intelligent agents in the marketplace) to 0.81 (when there are 30 such agents). This shows that by adding even a small number of intelligent agents to the marketplace the efficiency can be improved¹⁴. In Figure 9, Smith's Alpha decreases as the number of intelligent agent increases. This indicates that the market is slowly converging to its equilibrium and becoming more efficient as more and more intelligent agents enter the marketplace.

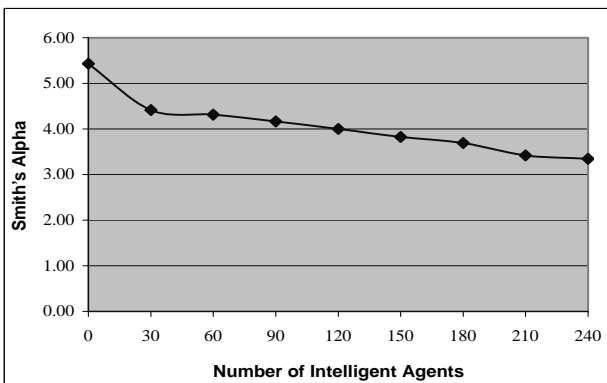


Figure 9: Relationship between Number of Intelligent Agents and Smith's Alpha

Finally, the relationship between the sellers' average utilities and the number of intelligent agents is shown in Figure 10.

¹⁴It should be pointed out here, that there are 240 intelligent agents in the marketplace and this makes up 80% of the total bidder population. It is predicted that the allocative efficiency would be much higher if the marketplace is completely populated with intelligent agents only. However, as described earlier, this is not possible because of the auctions' limitations.

Here, the average utility obtained by the sellers increases as the number of intelligent agents increases. When there are between 30 and 120 agents, the increase in average utility is minimal; but when the number is between 120 and 180, there is a significant increase in the seller's average utility. When there is a small number of intelligent agents (≤ 120), they compete with the dummy bidders and they gain a high utility as reflected in Table 3. However, when there are more than 120 agents in the marketplace, the agents are now competing amongst themselves resulting in a lower average utility but a higher gain to the sellers. In this case, it is more beneficial for the sellers to trade their goods in an environment where there are many bidders that adopt our intelligent strategy.

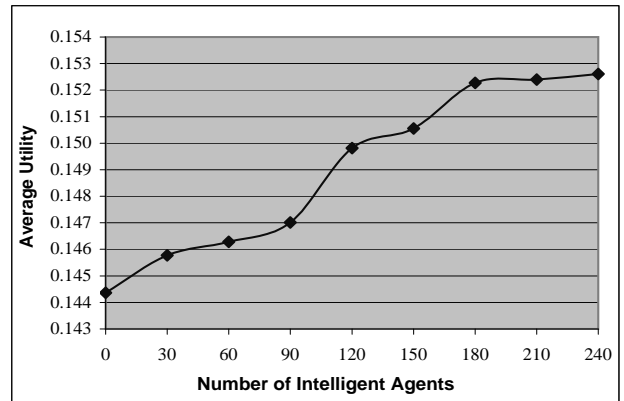


Figure 10: Relationship between Number of Intelligent Agents and Sellers' Average Utility

These results confirm our motivating hypothesis that our intelligent bidding strategy can perform effectively in a wide range of environments. It achieved a high success rate and average utility in almost all the cases considered and performed reasonably well even in cases where it cannot accurately assess the environment. Moreover, the agents achieved high utility even when there are many of them in the marketplace. We can also conclude that as the number of intelligent agents increases, the market efficiency increases which, in turn, leads to an increase in profit to the sellers.

4. RELATED WORK

There have been several attempts to design sophisticated and efficient bidding strategies for agents participating in online auctions. To this end, Faratin et al.'s model [6] is broadly similar to the one defined in this paper. However, there are several important differences between one-to-one negotiations and multiple auctions. Chief amongst these are the type of the tactics that are considered relevant and the aspects of the domain that need to be reflected in these tactics. An extension to Faratin's model is given by Matos et al. [7] who analysed the evolution of negotiation strategies using GAs, and determined which of them are appropriate in which situations. This approach is somewhat similar to our work, but the main difference is in the domain that we are dealing with (multiple auctions versus bi-lateral negotiations).

Preist et al. [10] proposed an algorithm design for agents that participate in multiple simultaneous English auctions. The algorithm proposes a coordination mechanism to be used in an environment where all the auctions terminate simultaneously, and a learning method to tackle auctions that terminate at different times. Byde [4] also considers this environment, but utilises stochastic dynamic programming to derive formal methods for optimal algorithm specification that can be used by an agent when participating in simultaneous auctions for a single private value good. Both of these works are designed specifically for purchasing items in multiple English auctions and their algorithms are not applicable in a heterogeneous protocol context. Byde et al. [5] presented another decision theoretic framework that an autonomous agent can use to bid effectively across multiple auctions with various protocols. The framework uses an approximation function that provides an estimate of the expected utility of participating in the set of future auctions and it can be employed to purchase single or multiple items. However, at this time the evaluation of the algorithm's operational effectiveness has not been reported and so we cannot determine whether it will outperform our heuristic method.

5. CONCLUSIONS AND FUTURE WORK

This paper presented a novel bidding algorithm that can be used for an agent to participate in multiple online auctions with varying start and end times and with multiple protocols. We extended the work presented in [2] and [3] and showed the strategies that were evolved for the complete set of our trading scenarios. These strategies were then codified into the agent's reasoning behaviour so that it can select the most appropriate strategy to employ in its prevailing circumstances. Through empirical evaluation we showed that our intelligent bidding strategy is effective in a wide range of environments. Moreover, this holds true even when there are many such agents in the system at the same time and when the agent cannot be completely accurate about the type of environment it is situated in. Also, as more of our agents entered the marketplace, the overall system becomes more competitive and efficient leading to an increase in the sellers' profit.

There are a number of areas of further investigation. We need to develop methods such that an agent can quickly and accurately determine the type of trading environment in which it is situated. While this is reasonably easy to do in our simulated environment, we believe it will be significantly more difficult in real world scenarios. Thus, although our method is robust to inaccurate environment assessment, it is better if this can be minimised. We would also like to compare our heuristic method with the more principled decision theoretic approach of [5] to determine the relative strengths and weaknesses of the two methods.

6. REFERENCES

- [1] P. Anthony. *Bidding Agents for Multiple Heterogeneous Online Auctions*. PhD thesis, University of Southampton, UK, 2003.
- [2] P. Anthony and N. R. Jennings. Evolving bidding strategies for multiple auctions. In *Proceedings of the 15th European Conference on Artificial Intelligence*, pages 178–182, 2002.
- [3] P. Anthony and N. R. Jennings. Developing a bidding agent for multiple heterogeneous auctions. *ACM Transaction on Internet Technology*, 3(3), 2003.
- [4] A. Byde. Programming model for algorithm design in simultaneous auctions. In *Proceedings of the 2nd International Workshop on Electronic Commerce*, pages 152–163, New York, USA, 2001. ACM Press.
- [5] A. Byde, C. Preist, and N. R. Jennings. Decision procedures for multiple auctions. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and MultiAgent Systems*, pages 613–620, New York, USA, 2002. ACM Press.
- [6] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
- [7] N. Matos, C. Sierra, and N. R. Jennings. Successful negotiation strategies: An evolutionary approach. In *Proceedings of the 3rd International Conference on Multi Agent Systems*, pages 182–189, Paris, France, 2001. IEEE Press.
- [8] R. P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:639–738, 1987.
- [9] M. Mitchell. *An introduction to genetic algorithms*. MIT Press, Cambridge, Mass., 1996.
- [10] C. Preist, C. Bartolini, and I. Phillips. Algorithm design for agents which participate in multiple simultaneous auctions. In *Agent Mediated Electronic Commerce III*, LNAI, pages 139–154, Berlin, German, 2001. Springer.
- [11] V. L. Smith. *Papers in Experimental Economics*. University Press, Cambridge, MA, 1992.