

Developing a Bidding Agent for Multiple Heterogeneous Auctions

PATRICIA ANTHONY and NICHOLAS R. JENNINGS

University of Southampton

Due to the proliferation of online auctions, there is an increasing need to monitor and bid in multiple auctions in order to procure the best deal for the desired good. 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 auctions with varying start and end times and with varying protocols (including English, Dutch and Vickrey). The framework is flexible, configurable, and enables the agent to adopt varying tactics and strategies that attempt to ensure that the desired item is delivered in a manner consistent with the user's preferences. Given this large space of possibilities, we employ a genetic algorithm to search (offline) for effective strategies in common classes of environment. The strategies that emerge from this evolution are then codified into the agent's reasoning behaviour so that it can select the most appropriate strategy to employ in its prevailing circumstances. The proposed framework has been implemented in a simulated marketplace environment and its effectiveness has been empirically demonstrated.

Categories and Subject Descriptors: I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*intelligent agents*; K.4.4 [**Computers and Society**]: Electronic Commerce—*distributed commercial transactions*

General Terms: Algorithms, Design, Experimentation

Additional Key Words and Phrases: multiple auctions, bidding strategy, genetic algorithms

1. INTRODUCTION

Online auctions are a popular and effective medium for procuring goods and services in both business to business and business to consumer electronic commerce [Bapna et al. 2001; He et al. 2003]. To date, there are more than 2500 auction houses that conduct business online.¹ Some of the well-known auction

¹<http://www.InternetAuctionList.com>.

P. Anthony also wishes to acknowledge the funding of University Malaysia Sabah.

Author's address: Intelligence, Agents, Multimedia Group, Department of Electronics and Computer Science, University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom; email: {pa99r, nrj}@ecs.soton.ac.uk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2003 ACM 1533-5399/03/0800-0185 \$5.00

houses include eBay,² Amazon.com,³ Yahoo!Auction,⁴ Priceline,⁵ UBid⁶ and many others. These auction houses conduct many different types of auctions, but the more popular ones are English, Dutch, first-price sealed bid and second-price sealed bid (also known as Vickrey). In an English auction, the auctioneer begins with the lowest acceptable price and bidders are free to raise their bids successively until there are no more offers to raise the bid or until the end of the auction is reached if there is a time limit. The winning bidder is the one with the highest bid [McAfee and McMillan 1987]. A Dutch auction is the opposite of an English one in that the auctioneer starts with an initial high price, which is then lowered progressively until there is an offer from a bidder to claim the item. In the first-price sealed bid, each bidder submits their offer for the item without any knowledge of the others bids. The highest bidder gets the item and they pay a price equal to their bid. The Vickrey auction is similar to the first-price sealed bid auction, where the item is awarded to the highest bidder but they pay a price equal to the second highest bid [Vickrey 1961].

As the number of auction houses increases, consumers are faced with the problem of monitoring many auction sites, picking which auction to participate in, and making the right bid to ensure that they get the desired item under conditions that are consistent with their preferences. These processes of monitoring, selecting and making bids are both complex and time consuming. The task becomes even more complicated when there are different start and end times and when the auctions employ different protocols. For this reason, some online auctions provide bidding agents (proxy bidders) to assist consumers with these tasks.⁷ In these systems, the consumer needs to supply their bidding agent with details of the item that they wish to purchase, as well as the maximum amount that they are willing to pay for it. The bidding agent then proceeds to bid in the auction and updates the consumer about the status of the auction from time to time until it terminates. While freeing the consumer to a certain extent, this type of facility typically has the problem of only being able to operate at a single auction site, or only operating with a single auction protocol. If the consumer wishes to participate in another auction, they need to wait until the auction is concluded to avoid getting multiple items. In addition, the auction site host could use information about the consumer to cheat them by manipulating the auction. This is possible in English auctions through shill bidding where bids are deliberately placed to drive up the price of an item [McAfee and McMillan 1992]. In this case, the consumer can be made to pay a price equal (or very close) to their maximum bid.

In view of these complexities, consumers tend to focus on a single auction of their choice. Unfortunately, winning in that auction does not necessarily mean

²<http://www.ebay.com>.

³<http://www.amazon.com>.

⁴<http://auctions.shopping.yahoo.com>.

⁵<http://www.priceline.com>.

⁶<http://www.ubid.com>.

⁷<http://www.ebay.com>.

that they get the best deal. They may have received a better price in another auction. The losers may also have had better luck in another auction as well. From the seller's perspective, the lock in of bidders is also a disadvantage since those agents that value the good highly may not be amongst this set. This, in turn, restricts the expected revenue from the auction.

As a step toward the multiple auctions case, consumers can utilize the services of auction search engines (such as BidXS,⁸ AuctionWatch,⁹ and AuctionBeagle¹⁰). These allow the consumer to monitor multiple concurrent auctions, but they leave the actual bidding decision to the consumer. While this certainly increases the consumer's knowledge of the global marketplace, it does not solve the problem of reducing the amount of time that has to be spent online. Moreover, deciding what amount to bid for an item requires an intelligent decision where the consumer needs to come up with a strategy to work out the bid value. In many cases, the outcome of this decision making is that the consumer is trapped with the *winner's curse* phenomenon where they pay more than the actual value of the item [Klemperer 1999].

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. 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 tries 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 article 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 Vickrey.¹¹ The strategies themselves are heuristic in nature because the multiple heterogeneous auction environment is very complex, dynamic and unpredictable, making it impossible to find an optimal strategy that can be used in practical contexts [Byde et al. 2002]. Moreover, our early investigations showed that the effectiveness of the strategies was heavily influenced by the nature of the environment [Anthony et al. 2001]. 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 [Anthony and Jennings 2002]. We chose this particular method because GAs have been known to perform well in areas where the space to be searched is large and not well understood

⁸<http://www.bidxs.com>.

⁹<http://www.auctionwatch.com>.

¹⁰<http://www.auctionbeagle.com>.

¹¹The first-price sealed bid is not considered here because of its similarities (from a strategic point of view) to the Dutch auction [McAfee and McMillan 1987].

[Mitchell 1996]. Having evolved the strategies, the agent adopts the one that is most appropriate to its prevailing context.

This paper advances the state of the art in the following ways. First, we develop a high level decision-making framework for an agent to bid across multiple concurrent auctions of varying protocols with varying start and end times. This framework is heuristic in nature and uses tactics and strategies to vary the agent's behavior so as to ensure a good fit with the user's auction objectives. Second, we evolve a strategy that is effective in our multiple auction context. This strategy consists of multiple evolved sub-behaviors that are appropriate in different environmental settings and with different user objectives. This strategy is termed the *intelligent bidding strategy* in the remainder of this article. Third, we systematically evaluate the intelligent bidding strategy to highlight its operational characteristics in different scenarios.

The remainder of the article is structured in the following manner: In the next section, we describe our electronic marketplace scenario in which our algorithms are evaluated. Section 3 describes the fundamentals of the bidding algorithm and our initial evaluation of the various environmental factors that effect its performance. Section 4 describes the work on evolving strategies for particular classes of environments and the evaluation of the effectiveness of this method of approach. Section 5 describes and evaluates the intelligent bidding strategy. Section 6 discusses related work and finally, Section 7 presents our conclusions and further work.

2. THE ELECTRONIC MARKETPLACE

The simulated electronic marketplace consists of a number of auctions that run concurrently (see Figure 1).¹² There are three types of auctions running in the environment: English, Dutch and Vickrey. The English and Vickrey auctions have a finite start time and duration generated randomly from a uniform probability distribution, the Dutch auction has a start time but no predetermined end time. The start and end times vary from one auction to another. At the start of each auction (irrespective of the type), a group of random bidders are generated to simulate other auction participants. These participants operate in a single auction and have the intention of buying the target item and possessing certain behavior. They maintain information about the item they wish to purchase, their private valuation of the item (the maximum amount that they are willing to pay for the desired item), the starting bid value and their bid increment. These values are generated randomly from a standard probability distribution. Their bidding behavior is determined based on the type of auction that they are participating in. In an English auction, they initiate bidding at their starting bid value; when making a counter offer, they add their

¹²The marketplace is a virtual simulation in that it is supposed to represent all the auctions that are selling the desired item anywhere on the Internet. It is obviously a simplification, since by grouping them in this way we can focus on the agent's bidding strategy (our main aim) and we do not have to worry about the (very important and difficult) problems of finding all the auctions that sell the desired item, semantic interoperability due to heterogeneous information models and ontologies, latency due to network traffic variability, or interfacing our software to proprietary commercial auction houses.

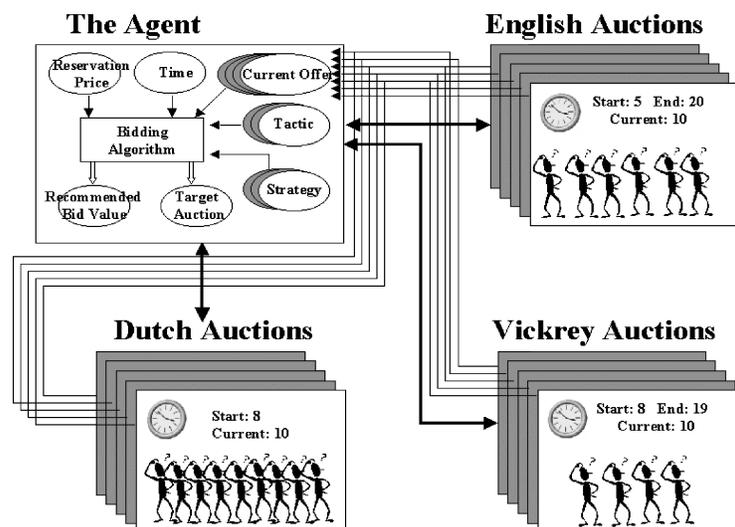


Fig. 1. The Marketplace simulator.

bid increment to the current offer (provided the total is less than their private valuation), and they stop bidding when they acquire the item or when their private valuation is reached. In a Dutch auction, they wait until the offer value is equal to or less than their private valuation before making an offer. Finally, in a Vickrey auction, they bid at their private valuation. These strategies are all based on the dominant strategies of the respective one-shot single auctions [Sandholm 1999].

The auction starts with a predefined starting value; a small value for an English auction and a high value for a Dutch one. There are obviously no start values for the Vickrey auctions. Offers and counter offers are accepted from bidders that are picked randomly from the pool of bidders in that particular auction. These processes are repeated until either the private valuation or the end time for that auction is reached. The winner in an English auction is the bidder with the highest bid value at the end of the auction. An item will remain unsold if the highest price offered is less than the item's reservation price. In a Dutch auction, when no offer is received from the bidders, the value is reduced (based on a fixed decrement value) and the whole process is repeated. The item is sold when a bidder agrees to buy the item at the current offer price. If there is more than one bidder who is interested at the same price, the item will be sold to the bidder who offered to buy the item first. There may be cases where there is no offer from any bidders throughout a particular auction. In this situation, the auction terminates when the decremented offer reaches the reservation price for the item. Bidders in a Vickrey auction submit their bid values before the end of the auction. Bids are opened at the end of the auction and the winner is the one who offered the highest price, but the amount paid for the item is the value of the second highest bid. If there is a tie, the winner is the bidder who submits the earliest bid. As in the English auction, an item will not be sold if the offer price is less than its reservation price.

The marketplace is flexible and can be configured to run any number of auctions and operate for any length of discrete time. We assume that all auctions in the marketplace are auctioning the item that the consumers are interested in. Our bidder agent is allowed to bid in any of the auctions and it aims to deliver the item to its consumer based on their preferences. The bidder agent is given a hard deadline by when it needs to obtain the item. The bidder agent utilizes the available information to make its bidding decision; this includes the consumer's private valuation, the time it has left to acquire the item, the current offer of each individual auction, and its sets of tactics and strategies. The private valuation is derived from the item's closing price distribution, observed from past auctions. The tactics and strategies are the main constituents that drive the agent's behavior in making the bidding decision (these are described in Section 3). The output of the bidding decision is the auction the agent should bid in and the recommended bid value that it should bid in that auction. If the agent does not purchase the item by its deadline, it returns to the consumer for further instructions.

3. DESIGNING THE AGENT'S BIDDING STRATEGY

To ensure that our bidder agent operates effectively in the marketplace, it needs to possess a strategy to ensure that it obtains the item within the given time according to the consumer's preferences. Here the bidding strategy for the agent is modelled on the idea of negotiation decision functions as proposed by Faratin et al. [1998]. The original model defined a range of strategies that an agent can employ to generate initial offers and counter offers in a two party negotiation. This model identifies the key constituents that drive an agent's negotiation behavior and defined a single tactic to deal with each of them. The agent's overall behavior is then the amalgamation of these different facets, weighted by their relative importance to the user. Mapping this to an auction environment, the bidder agent needs to decide the new bid value based on the current offer price. The current offer can be treated as the offer and the new bid value can be treated as the counter offer. Negotiation is over when the auction terminates or when the bidder's private valuation is reached (bidder walks out of the negotiation process).

First, we will present our notation. Let t be the current universal time across all auctions, where $t \in \Gamma$, and Γ is a set of finite time intervals. Let t_{\max} be the time by when the agent must obtain the good (i.e. $0 \leq t \leq t_{\max}$), and let A be the list of all the auctions that will be active before t_{\max} . At any time t , there is a set of active auctions $L(t)$ where $L(t) \subset A$. Let $E(t)$, $D(t)$, and $V(t)$ be the set of active English, Dutch and Vickrey auctions, respectively, where $E(t) \cap D(t) = \phi$, $D(t) \cap V(t) = \phi$, $E(t) \cap V(t) = \phi$, and $E(t) \cup D(t) \cup V(t) = L(t)$. Each auction $i \in A$, will have its own start time, σ_i , and its own end time η_i where $i \in E(t) \cup V(t)$. Let λ be the agent's bid increment value, and p_r be its private valuation for the target item. Let *Item_NA* be a Boolean flag to indicate whether the target item has already been purchased by the agent. We assume that the value of p_r is based on current reliable market prices observed from past auctions, and that the marketplace is offering the item that the agent is interested in.

```

while ( $t \leq t_{max}$ ) and ( $Item\_NA = true$ )
  Build active auction list;
  Calculate current maximum bid using the agent's strategy;
  Select potential auctions to bid in, from 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

```

Fig. 2. Top-level algorithm for the bidding agent.

With these definitions in place, the algorithm for the bidding agent is summarized in Figure 2. Since each auction has a different start and end time, the bidder agent needs to build an active auction list to keep track of all the auctions that are currently active in the marketplace. We define an active auction as one that has started but not yet reached its end time. The agent identifies all the active auctions and gathers relevant information about them. It then calculates the maximum bid it is willing to make at *the current time* using the agent's strategy (described in Section 3.1). This current maximum bid, by definition, will always be less than or equal to its private valuation. 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 (described in Section 3.2). The auction and corresponding bid with the highest expected utility is selected from the potential auctions as the target auction. Finally, the agent bids in the target auction.

3.1 Calculating the Current Maximum Bid

At any given time t , the agent needs to determine its current maximum bid. This bid is defined as the maximum value that the agent is willing to offer at *the current moment in time*. There are several factors that the agent needs to take into consideration when calculating this value. One is the remaining time that it has left to acquire the item. Thus, a key determinant of what price to offer depends on how much time it has left to bid. For example, the agent may decide to make a low bid value when it has a lot of time left, and as the remaining time decreases, the agent bids closer to its private valuation. Another consideration is the number of remaining auctions that the agent can bid in. Here the agent's behavior is similar to the time constraint, whereby it may choose to maximize its chances of success by bidding close to the private valuation when the number of auctions is small. The level of desperateness to obtain the item is another consideration that the agent may need to take into account. If the agent is desperate to get the item, it bids aggressively as soon as it starts to ensure that it maximizes its chance of getting the item. The opposite of this behavior is that of an agent that is looking for a bargain. If an agent wishes to get the item at a bargain price, it starts bidding at a very low price and eventually bids its private valuation when it has very little time left. The set of considerations of the remaining time left, the remaining auctions left, the desire to get a bargain and the level of desperateness are here referred to as the *bidding constraints*.

The agent uses some combination of these constraints in order to determine its maximum bid value at the current moment in time. Our model is open in that if there was another aspect that the consumer wanted the agent to consider, then this could easily be added in as a new bidding constraint. Exactly which constraints are used in a given situation is determined by the consumers and their preferences. Thus, a consumer who wants to ensure it receives the item as quickly as possible would place the greatest weight on the time until deadline and the desperateness tactics, whereas a consumer who is looking to minimise the price paid would value the bargain tactic most highly.

More formally, let C be the set of considerations that the agent takes into account when formulating a bid and j represent the individual bidding constraints, where $j \in 1, \dots, |C|$. Let Δt denote the remaining time left for the agent to bid (i.e., $t_{\max} - t$) and Δa denote the number of auctions left in the marketplace. Let μ denote the agent's desire for a bargain, where $\mu \in [0, 1]$ (where 1 is actively looking for a bargain and 0 is not actively looking for a bargain), and ε denotes the agent's level of desperateness for the item where $\varepsilon \in [0, 1]$ (where 1 is very desperate and 0 is not desperate). For each of constraint $j \in C$, there is a corresponding function $f_j(t)$ which suggests the value to bid based on that constraint. These individual constraints are then combined using a function F to produce the agent's overall position. Examples for F include weighted average, max or min.

At a 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 point in time). In this work, if the agent combines multiple bidding constraints, it allocates weights to denote their relative importance. Thus, let $w_j(t)$ be the weight on constraint j at time t , where $\forall j \in C, 0 \leq w_j(t) \leq 1$, and $\sum_{j \in C} w_j(t) = 1$.

The current maximum bid value for the agent at time t , is then calculated as $M(t) = \sum_{j \in C} w_j(t) f_j(t)$. The agent uses a set of polynomial functions (drawn from Faratin et al.'s [1998] negotiation functions) to calculate the bid value based on a single bidding constraint. Here this set of functions is referred to as the *tactics*. In the current implementation, the four tactics are remaining time, remaining auctions, desire for bargain and desperateness. The definition of each of these is given below.

3.1.1 The Remaining Time Tactic. This tactic determines the recommended bid value based on the amount of time remaining for the agent. Assume that the agent is bidding at time $0 \leq t \leq t_{\max}$. The agent bids closer to p_r as t approaches t_{\max} , and it eventually bids at p_r when $t = t_{\max}$. To calculate the bid value at time t , the following expression is used:

$$f_{rt} = \alpha_{rt}(t) p_r$$

where $\alpha_{rt}(t)$ is a polynomial function of the form:

$$\alpha_{rt}(t) = k_{rt} + (1 - k_{rt}) \left(\frac{t}{t_{\max}} \right)^{\frac{1}{p}} \quad (1)$$

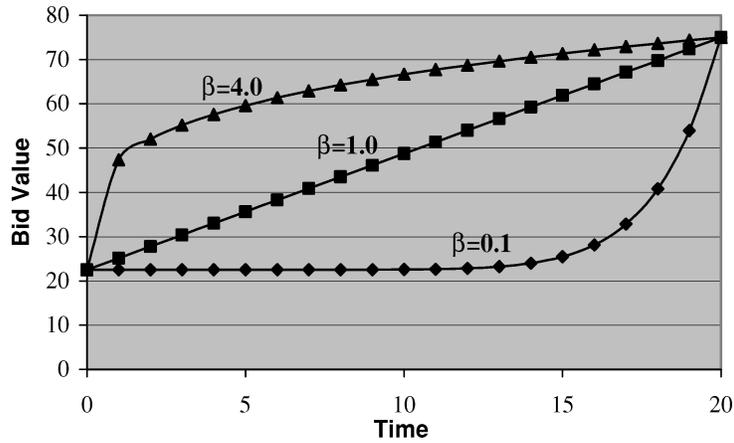


Fig. 3. Bid value for remaining time tactic with $t_{\max} = 20$ and $p_r = 75$.

k_{rt} is a constant that when multiplied by the size of the interval determines the value of the starting bid of the agent in any auction. By varying the value $\alpha_{rt}(t)$, a wide range of time-dependent functions can be defined from those that start bidding near p_r quickly, to those that only bid near p_r right at the end, to all possibilities in between. The only conditions are that $0 \leq \alpha_{rt}(t) \leq 1$, $\alpha_{rt}(0) = k_{rt}$, $\alpha_{rt}(t_{\max}) = 1$, and $0 \leq k_{rt} \leq 1$.

Using the polynomial function defined in Eq. (1), different shapes of curve can be obtained by varying the values of β . This represents an infinite number of possible tactics, one for each value of β . In this tactic, β is drawn from \mathfrak{R}^+ , where $0.005 \leq \beta \leq 1000$. When $\beta < 1$, the tactic maintains a low bid value until the deadline is almost reached, where this tactic concedes by suggesting the private valuation as the recommended bid value. The other extreme is when $\beta > 1$; the tactic starts with a bid value close to the private valuation and quickly reaches the reservation value long before the deadline is reached. By way of illustration, Figure 3 shows the different convexity degrees of the curves with $k_{rt} = 0.30$ and β taking on the values 0.1, 1.0 and 4.0, respectively.

3.1.2 The Remaining Auctions Tactic. This tactic is broadly similar to the remaining time tactic; the agent bids closer to p_r as the number of remaining auctions approaches 0 (since it is running out of opportunities to purchase the desired item). Thus, f_{ra} has the same form as f_{rt} and α_{ra} is defined as follows:

$$\alpha_{ra} = k_{ra} + (1 - k_{ra}) \left(\frac{c(t)}{|A|} \right)^{\frac{1}{\beta}}$$

Most of these terms are similar to α_{rt} , the only difference being that $c(t)$ is the list of auctions that have been completed between time 0 and time t . β is again drawn from \mathfrak{R}^+ , where $0.005 \leq \beta \leq 1000$.

3.1.3 The Desire for Bargain Tactic. This tactic is employed when the agent is motivated to try and obtain a bargain. The agent keeps the λ to a minimum as

it progresses from $t = 0$ to t_{\max} , but eventually bids its private valuation when t_{\max} is reached. To determine the bid value for this tactic, the agent considers the minimum bid value for the target item across all the auctions in the marketplace. At a given time t , newly started English auctions have low current bid values and Dutch auctions have very high current bid values. On the other hand, when auctions are terminating, English auctions typically have high current bid values and Dutch auctions have low current bid values. Vickrey auctions do not have any information on the bid values since bids are sealed and they are only opened at the end time. To deal with these points, the minimum bid value is calculated by taking into consideration the current bid value and the proportion of time left in the auction. These values are summed and averaged with respect to the number of active auctions at that time.

Let $v_i(t)$ be the current highest bid value in an auction i at time t , where $i \in L(t)$, and $\omega(t)$ be the minimum bid value for the agent at time t where:

$$\omega(t) = \frac{1}{|L(t)|} \left(\sum_{1 \leq i \leq |L(t)|} \frac{t - \sigma_i}{\eta_i - \sigma_i} v_i(t) \right).$$

The bid value is then calculated using the expression:

$$f_{ba} = \omega(t) + \alpha_{ba}(t)(p_r - \omega(t))$$

where $\alpha_{ba}(t)$ is defined as:

$$\alpha_{ba}(t) = k_{ba} + (1 - k_{ba}) \left(\frac{t}{t_{\max}} \right)^{\frac{1}{\beta}}. \quad (2)$$

Assume that $\alpha_{ba}(t)$ is similar to the polynomial function discussed in the first two tactics, but this time, $0.1 \leq k_{ba} \leq 0.3$, the minimum value of β equals 0.005 and the maximum value of β equals 0.5. These values reflect the fact that an agent that is looking for a bargain should never bid with $\beta > 1$ because this would inflate the agent's bid well before the deadline. In contrast, an agent that is looking for a bargain (with $\beta < 1$) maintains a low bid value until the deadline is almost reached where it will then suggest p_r as the recommended bid value. By conceding with a recommended bid value of p_r , the agent tries to ensure that it still successfully acquires the item even if it did not succeed in getting a bargain.

3.1.4 The Desperateness Tactic. This tactic is employed when the agent is desperate to get the item. The agent bids close to p_r at $t = 0$, and eventually bids at p_r when t_{\max} is reached. In this tactic, the agent utilizes the minimum bid value and the polynomial function of Eq. (2) but with a slight variation to the value of β , where $1.67 \leq \beta \leq 1000$ and $0.7 \leq k_{de} \leq 0.9$. The values picked for k_{de} are high since a desperate agent starts bidding at a value that is near p_r . With these minor variations, f_{de} is the same as f_{ba} and α_{de} is the same as α_{ba} .

3.2 Selecting Potential Auctions and the Target Auction

The agent selects potential auctions if and only if it is not holding the highest bid in an English auction or it has not placed a bid in a Dutch or a Vickrey auction. This is to ensure that the agent does not acquire more than one target item. The agent selects the potential auctions by considering values for the current maximum bid for each active auction. In the English auctions, this is carried out by taking those auctions that are close to their end time, in which the current bid value when added to the bid increment is less than or equal to the current maximum bid. The agent's new bid value is the current bid plus the bid increment. Only English auctions that are close to their end time are picked to maximise the agent's chances of winning. If the agent currently holds a bid in an English auction that still has a long time to complete, it will not be able to participate in other auctions until it loses out to another bidder or until the auction terminates. There are several potential outcomes when an auction terminates; the agent loses out to another bidder; the agent's bid value may be less than the private valuation (in which case there will be no winner); or the agent wins. If either of the first two situations occur, the agent loses out in that it wasted a lot of the time in one auction thus reducing its chances of participating in other auctions. The potential Dutch auctions in which the agent may bid are those with current bids that are less than the current maximum bid. Here, the agent's new offer is the current bid for that particular Dutch auction. The potential Vickrey auctions in which the agent may bid are those that are terminating at the current time and the agent's bid value is its current maximum bid value. The reason for this is because the timing of the bid does not affect the outcome of the auction, and so it is better to wait until the last minute [Byde et al. 2002] before placing a bid in a given Vickrey auction.

If there is only one auction in the potential auction list, that one is picked as the target auction. If there are multiple auctions, the agent calculates the expected utility for each of them. By definition, the expected utility is the product of the probability of the agent winning in that auction at the given bid value and the value of the agent's utility function. The auction with the highest expected utility for the agent's bid value is picked as the target auction. To calculate the agent's probability of winning in an auction, we take into account the closing price distribution for all the auctions observed over time. In more detail, let $P_i^w(v)$ be the agent's probability of winning in auction i if it bids with the value v . Let $P_i^c(p)$ be the probability of auction i closing at price p . The probability of winning in an English or Vickrey auction is then:

$$P_i^w(v) = \sum_{p>v} P_i^c(p) + \frac{P_i^c(v)}{2}$$

where

$$\sum_{p<v} P_i^c(p) \leq P_i^w(v) \leq \sum_{p \leq v} P_i^c(p).$$

We assume that there is a possibility that more than one bidder may bid in an auction i at value v . The probability of winning in a Dutch auction is slightly different from the probability of winning in English and Vickrey auctions, due

to the nature of the auction. In this situation, when the bid value is at p_r , it is almost certain that if our agent bids at that value it will win the auction unless there are collisions (when there is more than one bidder bidding for that item at a particular price). Let n_i be the number of bidders offering to buy the item at the current offer price in Dutch auction i , and let x be the number of Dutch auctions. The average number of collisions, θ , is then calculated as

$$\theta = \frac{\sum n_i}{x}.$$

Hence, the probability of winning in a Dutch auction is defined as $P_i^w(v) = \frac{1}{\theta}$. The expected utility for an auction i with a bid value v is then calculated as:

$$u_i(v) = P_i^w(v)U_i(v),$$

where

$$U_i(v) = 1 - \left(\frac{v}{p_r}\right)^{\frac{1}{\beta}}.$$

The β used here is the same as the one used in generating the polynomials for the tactics. The utility function for each potential auction is calculated by dividing the payoff amount with p_r . The utility value is higher when the payoff ($p_r - v$) is high (value greater than 0), and it is lower when the payoff is low (value close to or equal to 0).

3.3 Evaluating the Bidding Strategy

To evaluate the performance of our agent using the bidding algorithm described above, we undertook an empirical evaluation. We ran the experiments in an environment where $t_{\max} = 20$. The total number of auctions running is between 20 and 30 and each auction has between 2 and 10 participants. The participants of the English and Vickrey auctions use the optimal (dominant) strategies to their respective one-shot auctions. For Dutch auctions, the participants wait until the offer value is just less than their p_r before making an offer. In all cases, the p_r of these agents is drawn from the same probability distribution as that of our bidding agent. Three control models are defined as a basis for comparisons. The first model (C1) simulates the behavior of a typical simple bidding agent that randomly joins one auction and stays there until its p_r has been reached or until the auction is over. In the second model (C2), the agent picks the auction that has the closest end time where the current bid is less than its p_r , and stays there until the auction is over or until its p_r is reached. The third model (C3) is similar to C2 but the target auction is selected randomly. In both C2 and C3, if the item is not acquired, the process is repeated until the allocated time is over.

Experiments were conducted separately for different values of the user's private valuation. Eight different experiments were conducted where each control model and bidding agent used p_r values of 68, 70, 72, 74, 75, 77, 80 and 82. The bidding strategies for the three bidding agents were chosen based on our intuition depending on the value of the private valuation. Table I shows the

Table I. Strategies for Bidding Agents

P_r	Agent	(k_{rt}, β_{rt})	(k_{ra}, β_{ra})	(k_{ba}, β_{ba})	(k_{de}, β_{de})	$(w_{rt}, w_{ra}, w_{ba}, w_{de})$
68	I	(0.6, 500)	(0.6, 500)	(0.3, 0.5)	(0.7, 500)	(0.4, 0.2, 0.1, 0.3)
	II	(0.8, 900)	(0.8, 900)		(0.8, 900)	(0.4, 0.2, 0.0, 0.4)
	III	(0.9, 950)			(0.9, 950)	(0.5, 0.0, 0.0, 0.5)
70	I	(0.6, 100)	(0.6, 100)	(0.3, 0.5)	(0.7, 200)	(0.4, 0.2, 0.1, 0.3)
	II	(0.7, 500)			(0.8, 700)	(0.5, 0.0, 0.0, 0.5)
	II	(0.8, 900)			(0.9, 900)	(0.5, 0.0, 0.0, 0.5)
72	I	(0.6, 500)			(0.7, 500)	(0.5, 0.0, 0.0, 0.5)
	II	(0.6, 100)	(0.6, 100)	(0.3, 0.5)	(0.7, 200)	(0.25, 0.25, 0.25, 0.25)
	III	(0.6, 100)	(0.6, 100)	(0.3, 0.5)	(0.7, 200)	(0.4, 0.2, 0.1, 0.3)
74	I	(0.6, 150)	(0.6, 150)	(0.3, 0.5)	(0.7, 150)	(0.4, 0.2, 0.1, 0.3)
	II	(0.6, 200)			(0.7, 200)	(0.5, 0.0, 0.0, 0.5)
	III	(0.4, 500)			(0.4, 500)	(0.5, 0.0, 0.0, 0.5)
75	I	(0.5, 100)	(0.5, 100)	(0.3, 0.5)	(0.5, 100)	(0.4, 0.2, 0.1, 0.3)
	II	(0.6, 100)	(0.6, 100)	(0.3, 0.5)	(0.7, 200)	(0.4, 0.2, 0.1, 0.3)
	III	(0.6, 300)			(0.7, 300)	(0.5, 0.0, 0.0, 0.5)
77	I	(0.6, 100)	(0.6, 100)	(0.3, 0.5)	(0.7, 200)	(0.25, 0.25, 0.25, 0.25)
	II	(0.5, 100)			(0.6, 50)	(0.5, 0.0, 0.0, 0.5)
	III	(0.5, 50)			(0.5, 50)	(0.5, 0.0, 0.0, 0.5)
80	I	(0.3, 50)	(0.3, 50)	(0.2, 0.5)	(0.3, 50)	(0.25, 0.25, 0.25, 0.25)
	II	(0.2, 100)			(0.2, 100)	(0.5, 0.0, 0.0, 0.5)
	III	(0.2, 100)		(0.3, 0.5)	(0.2, 100)	(0.4, 0.0, 0.2, 0.4)
82	I	(0.1, 20)				(1.0, 0.0, 0.0, 0.0)
	II	(0.1, 20)		(0.3, 0.5)		(0.3, 0.0, 0.7, 0.0)
	III	(0.1, 20)	(0.1, 20)	(0.2, 0.5)	(0.1, 20)	(0.25, 0.25, 0.25, 0.25)

different strategies that were employed by the bidding agents for each private valuation value.

As an example, Agent I with $p_r = 68$, values remaining time as the most important aspect ($w_{rt} = 0.4$), followed by the desperateness behavior ($w_{de} = 0.3$), remaining auctions left ($w_{ra} = 0.2$) and bargain ($w_{ba} = 0.1$). It has a high starting bid value and quickly reaches p_r for the remaining time tactic and the remaining auctions left tactic. Its desire for a bargain is low ($k_{ba} = 0.3$, $\beta_{ba} = 0.5$) but its level of desperateness is very high ($k_{de} = 0.7$, $\beta_{de} = 500$). Agent II with $p_r = 68$ considers the remaining time left tactic and desperateness tactic as equally important ($w_{rt} = 0.4$, $w_{de} = 0.4$). It considers the remaining auctions left tactic as less important ($w_{ra} = 0.2$) and is not looking for a bargain at all ($w_{ba} = 0.0$). It starts with a high bid value ($k_{rt} = 0.8$, $\beta_{rt} = 900$, $k_{ra} = 0.8$, $\beta_{ra} = 900$, $k_{de} = 0.8$, and $\beta_{de} = 900$) and quickly reaches p_r for all the three tactics considered (this is why the column for the bargain tactic is left blank). The last agent for $p_r = 68$ (Agent III), considers the remaining time left tactic and desperateness tactic as very important ($w_{rt} = 0.5$, $w_{de} = 0.5$), and it is not interested in the remaining auctions left tactic or the bargain tactic. For both tactics ($k_{rt} = 0.9$, $\beta_{rt} = 950$, $k_{de} = 0.9$, and $\beta_{de} = 950$), the agent starts with a high bid value and quickly reaches p_r . These strategies are chosen because the prior probability of the agent winning in an auction is very small (0.06). As the probability of winning in an auction increases (by increasing its private valuation), the behavior of each tactic is changed accordingly as can be seen from Table I.

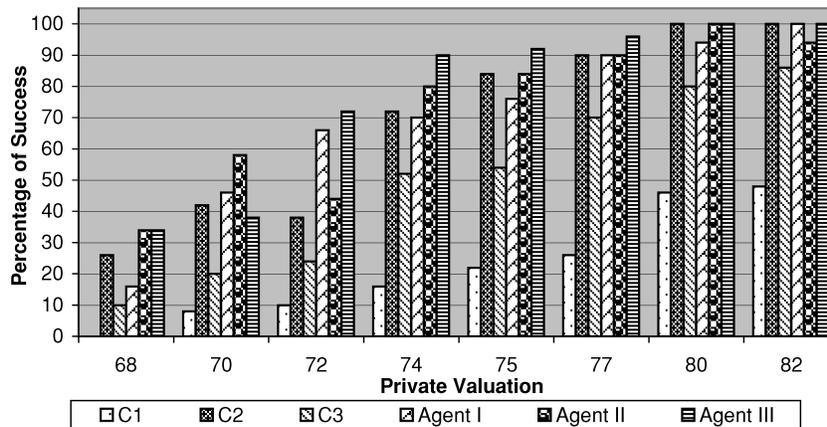


Fig. 4. Success rate comparisons.

For each p_r , the three bidding agents with varying strategies and the three control models were run for 100 times in the same environment, independently. Analysis of variance (ANOVA) was used to test the hypothesis about the differences between the success rate means (when running the experiment 50, 100, 150 and 200 times) [Cohen 1995]. The procedure revealed that for all the experiments, the differences between means were not significant ($F_{3,8} = 0.027$, $p > 0.05$) and so the results obtained are statistically significant.

Figure 4 shows the performance of the agents and the control models in terms of the success rate percentage. The success rate is the number of times, as a percentage, the agent is successful in obtaining the item. The first three bars represent the control models and the remaining three bars represent the three agents with three different bidding strategies. It can be seen that when p_r is low, the success rate is low. When p_r is high, the agents have a better chance of obtaining the desired item. At each p_r point, the bidding agents performed better than the control models. When the value of p_r reaches 80 and 82, C2, Agent I, Agent II, and Agent III achieved a 100% success rate, since the prior probability of winning in a single auction at this p_r is very close to 1. In this particular implementation, it is observed from the closing price distribution that the mean is 76 and the standard deviation is 5. This leads to the assumption that if the agent's p_r is 76, the probability of it winning in a single auction is 0.5.

The comparison for the average payoffs is shown in Figure 5. The average payoff is defined as $\sum_{1 \leq n \leq 100} \frac{(p_r - v_i)}{n}$, where v_i is the winning bid value for auction i . The average payoff is calculated by deducting the agent's bid value (the value at which it acquires the item) from the agent's private valuation. This value is then divided by the agent's private valuation, summed and averaged over the number of runs (in this case 100). Here, all three agents outperformed the control models for all p_r values. The payoffs obtained by the control models are consistently lower than the bidding agents' payoffs. It can also be seen that as p_r increases, the payoffs that the bidding agents received also increased. This indicates that the bidding agents actively looked for a bargain when their p_r is high.

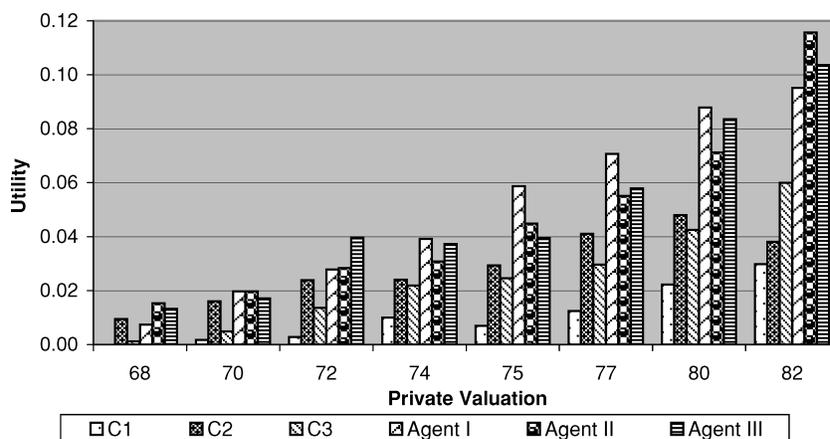


Fig. 5. Average payoff comparisons

The results of these experiments led to several conclusions. First, an agent that can bid effectively across multiple auctions is better than one that stays situated in a single auction. By being able to participate in multiple auctions, the agent is able to get a bargain (if it wishes to) by just changing the values of the bidding constraints. Secondly, the private valuation is one of the most important factors that needs to be considered when determining the strategy that should be employed by the agent. This is important, for example, since an agent with a very low p_r cannot look for a bargain and the agent should therefore consider this when accepting the user's preferences. The third observation is that the remaining time and remaining auctions tactics are the key determinants of successful behavior. Finally, the strategies to be used by the agent need to be dynamic, since not all strategies work well in all situations. Thus, a successful strategy in one situation may perform badly in another. However, it is possible to determine that certain classes of strategy are effective in environments that have particular characteristics. As an example, the agent strategy with values $k_{rt} = 0.60$, $\beta_{rt} = 100$, $k_{ra} = 0.60$, $\beta_{ra} = 100$, $k_{ba} = 0.30$, $\beta_{ba} = 0.50$, $k_{de} = 0.70$, $\beta_{de} = 200$, $w_{rt} = 0.40$, $w_{ra} = 0.20$, $w_{ba} = 0.10$ and $w_{de} = 0.30$ was used three times for p_r values of 70, 72 and 75 respectively. It can be seen that it achieved the highest success rate when p_r is 72 but placed second highest when the private valuation is at 70 and 75. Its average payoff is highest at $p_r = 72$ and $p_r = 70$ but it achieved second highest when $p_r = 75$. Here, we can conclude that this strategy performs best at $p_r = 72$, but it does not necessarily perform well with other values of p_r . However, this strategy did not perform well at $p_r = 75$ (even though this value is higher than 72) because the allocated bidding time may be short and the number of auctions that it can participate in may be limited. In this case, the key defining characteristics of an environment were found to be the number of auctions that are active before t_{\max} and the time the agent has to purchase the item (i.e., $t_{\max} - t_{\text{current}}$).

Having shown the benefits of having an agent that acts across multiple auctions and the close link between strategy effectiveness and environment type,

0.6	500	0.6	500	0.3	0.5	0.7	500	0.4	0.2	0.1	0.3
k_{rt}	β_{rt}	k_{ra}	β_{ra}	k_{ba}	β_{ba}	k_{de}	β_{de}	w_{rt}	w_{ra}	w_{ba}	w_{de}

Fig. 6. Encoding of a bidding strategy.

the next step is to determine which strategies should be used in which environments. To this end, the next section discusses how we used GAs to address this problem.

4. EVOLVING BIDDING 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 the given tactics and the weights for each tactic when these are combined. The number of strategies that can be employed is infinite, so, therefore, is the search space. Thus, handcrafting strategies, as per the previous section, is not realistic in the long term. Thus a means of automating the process of finding successful strategies is necessary. For the reasons outlined previously, we decided to use GAs to search offline for the most successful strategies in the predefined environments. Based on the results of Section 3.3, the three key determinants for the strategy selection are the remaining time left, the remaining auctions left and the private valuation. In this work, we defined four environments that take into accounts these determinants. The first one (STLA) is where there is a short bidding time ($10 \leq t_{\max} \leq 20$) and a small number of active auctions in the marketplace ($|L(t)| \leq 10$). The second environment (STMA) is where there is a short bidding time but the number of active auctions is large ($11 \leq |L(t)| \leq 45$). The third environment (LTLA) is where the allocated bidding time is long ($21 \leq t_{\max} \leq 100$) and where the number of active auctions is small. Finally, the last environment (LTMA) is where there is a long bidding time with many active auctions in the marketplace. Naturally, finer subdivisions are possible (see Section 5) but the focus here is demonstrating that strategies can be successfully evolved for broad classes of environments.

4.1 Encoding the Strategies

The individuals in the populations are the bidding agents and their genes consist of the parameters of the four different tactics and the relative weight for each tactic. Thus, the individuals are represented as an array of floating points values of:

- (1) k and β for the remaining time left tactic
- (2) k and β for the remaining auctions left tactic
- (3) k and β for the desire for a bargain tactic
- (4) k and β for the desperateness tactic
- (5) the relative weights for the four tactics.

By means of illustration, Figure 6 exemplifies the representation of the bidding strategy for Agent I with $p_r = 68$ in Table I. Each of the entries in the

column is treated as the gene. The first two columns indicate the agent's values for k and β for the remaining time left tactic, the next columns are the values for k and β for the remaining three tactics and the last four columns represent the relative weight for each tactic.

4.2 Computing the Fitness Function

The fitness function measures how well the individual performs against the others. Designing the fitness function is one of the key facets of GAs and so here we consider three plausible alternatives. These are the individual success rate in obtaining the item (Fitness Equation 1) and two variations based on the average utility. In the first case (Fitness Equation 2), the agent gets a utility of 0 if it fails to obtain the item. If it is successful, 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 ranging from 0.001 to 0.005 to ensure that the agent receives some value when the winning bid is equivalent to its private valuation. The final utility function (Fitness Equation 3) is similar to Fitness Equation 2 but the individual is penalised if it fails to get the item. In this case the penalty incurred ranges from 0.01 to 0.05. These values were chosen to analyse how the population evolves with varying degrees of penalty. Intuitively, Fitness Equation 1 should be used if delivery of the item is of utmost important, Fitness Equation 2 should be used if the agent is looking for a bargain and Fitness Equation 3 should be used when delivery of the item and looking for a bargain are equally important. The fitness score is then computed by taking the average utility from a total of 2000 runs. It is necessary to run these 2000 times to decrease the estimated standard error of the mean to a minimal level (when the number of runs is 500, the standard error of the mean is 5.0458, but this figure is significantly reduced to 1.1559 when the number of runs is increased to 2000). Analysis of Variance (ANOVA) was also used to test hypothesis about the differences between the fitness score means collected for the different number of runs. The null hypothesis of equal means was rejected because the procedure revealed that for all the experiments, the differences between means were significant ($F_{4,245} = 675.182$, $p < 0.05$) and so the results are statistically significant.

4.3 Searching for Successful Strategies

The algorithm for searching for acceptable strategies in a given environment is shown in Figure 7 and is elaborated upon in the remainder of this subsection.

4.3.1 Create Initial Bidder Population. The initial bidder population represents the starting point of the search and consists of $N = 50$ individuals that are generated randomly from the range of specified values. These values are based on the polynomial functions defined in Section 3.1. For the remaining time and the remaining auctions left tactics, the values are $0 \leq k \leq 1$ and $0.005 \leq \beta \leq 0.5$. In the desire for a bargain tactic, the values are $0.1 \leq k \leq 0.3$ and $0.005 \leq \beta \leq 0.5$. The values for k and β in the desperation tactic are $0.7 \leq k \leq 0.9$ and $1.67 \leq \beta \leq 1000$.

```

Randomly create initial bidder populations;
While not (Stopping Criterion) do
  Calculate fitness of each individual by running the
  marketplace 2000 times;
  Create new population
    Select the fittest individuals (HP);
    Create mating pool for the remaining
    population;
    Perform crossover and mutation in the
    mating pool to create new generation(SF);
    New generation is HP + SF;
End while

```

Fig. 7. The strategy searching genetic algorithm.

4.3.2 The Selection Process. The purpose of the selection process is to ensure that the fitter individuals are chosen to populate the next generation in the hope that their offspring will in turn have higher fitness [Beasley et al. 1993a]. “Elitism” is used here to force the GAs to retain some number of the best individuals at each generation [Mitchell 1996], since such individuals can be lost if they are not chosen to reproduce or if they are destroyed by crossover and mutation. Ten percent of the best individuals are copied to the new population to ensure that a significant proportion of the fitter individuals make it to the next generation. The remaining ninety percent of the individuals in the population are then chosen using Tournament Selection [Blickle and Thiele 1995b]. The selection is performed by choosing some number φ from the population and the best individual in this group is copied into the intermediate population (which is referred to as the mating pool). This process is repeated for 90% of N times. This selection technique is known to work well since it allows a diverse range of fit agents to populate the mating pool [Blickle and Thiele 1995a]. Once the mating pool is created, the individual with highest fitness is selected and moved to the new generation. The remaining individuals go through the process of crossover and mutation before making it to the new population. The new population includes a group of the fittest individual and the offspring generated from the reproduction process.

4.3.3 The Crossover Process. This process exchanges the genes between individual agents. Two individuals are randomly selected from the mating pool with crossover probability of $p_c = 0.6$, and the crossover point (c) is equal to 2. Crossover probability is the rate at which the population exchanges genetic materials [Mitchell 1996]. More specifically, two individuals are picked from the population. Two crossover points are then randomly picked. These points are where the two individuals will exchange their genetic material. The exchanging of genetic material process is performed using an extension combination operator [Beasley et al. 1993b], which works by taking the difference between two values of the crossover point, adding this difference to the higher (giving the maximum range) and subtracting it from the lower (giving a minimum range).

The new values are then generated between the minimum and maximum range.

4.3.4 The Mutation Process. Mutation allows the population to explore the search space but at a slower rate. In this work, individuals from the population are selected to mutate with a probability of $p_m = 0.02$. The gene from the chosen individual is picked randomly and a small value (0.05) is added or subtracted, depending on the range limitation for that particular gene. The mutation process is only applied to the values of k and β for each tactic. The weights are not considered here because adding a small value to the weight requires a renormalization and will have very little effect on the agent's online behavior.

4.3.5 The Stopping Criterion. The process terminates when the population converges. This is a condition where the population evolves over successive generations such that the fitness of the best and the average individual in each generation increase toward a global optimum [Beasley et al. 1993a]. The global optimum is defined as the highest peak in the search space [Mitchell 1996]. In this case, the population always converges before 50 iterations (typically, the value lies between 24 and 40).

4.4 Evaluating the Evolved Strategies

The aim of these experiments is to show that GAs can be used to evolve strategies that are effective in particular environments. To this end, the GAs are run in four different environments (in which the agent's p_r is set to 75). For each environment, we use the three different fitness functions described in Section 4.2. Apart from determining the strategies that work well in a given context, these experiments also aim to evaluate the strategies in terms of their success rate and the average payoff in a similar manner to the experimental evaluation conducted in Section 3.3. However, the key difference is that, here, the performance of the agents is evaluated based on an environment that has a particular set of characteristics. The performance of the evolved strategies is then compared with that of a control model C . C 's strategy is to bid in the auction that has the closest end time where the current bid is less than its p_r . This model was chosen because it performed well in the previous experiment as reported in Section 3.3 (in fact, this is C2). We also ran another set of experiments in the subenvironment of short time less auctions (STLA) in which the value of p_r is varied between a low value of 68, a medium value of 76 and a high value of 82. The purpose of this is to determine how the strategies evolve when varying p_r .

Turning to the first set of experiments (summarized in Table II). These results show the best strategies that have evolved for the different classes of environment. Each row contains the resulting strategies for each environment using Fitness Equations 1, 2 and 3.¹³ The values for the tactics are expressed as a pair of k and β and the weights for the bidding constraints are expressed as $(w_{rt}, w_{ra}, w_{ba}, w_{de})$. When a particular tactic is not present in the evolved

¹³FE1 indicates Fitness Equation 1, FE2 indicates Fitness Equation 2 and FE3 indicates Fitness Equation 3.

Table II. Summary of the Best Strategies with $p_r = 75$

Environment	(k_{rt}, β_{rt})	(k_{ra}, β_{ra})	(k_{ba}, β_{ba})	(k_{de}, β_{de})	$(w_{rt}, w_{ra}, w_{ba}, w_{de})$
FE1STLA	(0.73, 99.59)			(0.84, 56.04)	(0.76, 0.00, 0.00, 0.24)
FE1STMA	(0.63, 515.67)		(0.28, 0.31)	(0.73, 385.75)	(0.45, 0.00, 0.01, 0.54)
FE1LTLA		(1.00, 0.36)		(0.83, 67.38)	(0.00, 0.46, 0.00, 0.54)
FE1LTMA	(0.23, 683.97)			(0.78, 2.70)	(0.83, 0.00, 0.00, 0.17)
FE2STLA	(0.89, 1.44)	(0.94, 233.50)	(0.15, 0.40)	(0.71, 55.44)	(0.35, 0.16, 0.15, 0.34)
FE2STMA	(10.59, 507.92)	(0.81, 6.31)	(0.23, 0.06)	(0.80, 68.07)	(0.17, 0.03, 0.25, 0.55)
FE2LTLA	(0.70, 8.28)	(0.40, 5.21)	(0.25, 0.32)	(0.83, 648.90)	(0.65, 0.21, 0.02, 0.12)
FE2LTMA	(0.81, 9.74)	(1.00, 0.83)	(0.23, 0.04)	(0.82, 575.00)	(0.14, 0.49, 0.22, 0.15)
FE3STLA		(0.72, 25.56)		(0.87, 55.75)	(0.00, 0.42, 0.00, 0.58)
FE3STMA	(0.83, 52.00)		(0.10, 0.29)		(0.80, 0.00, 0.20, 0.00)
FE3LTLA	(0.41, 720.61)	(0.27, 7.95)		(0.71, 9.12)	(0.57, 0.19, 0.00, 0.24)
FE3LTMA	(0.59, 0.33)	(0.21, 654.55)	(0.12, 0.03)	(0.89, 19.17)	(0.16, 0.05, 0.01, 0.78)

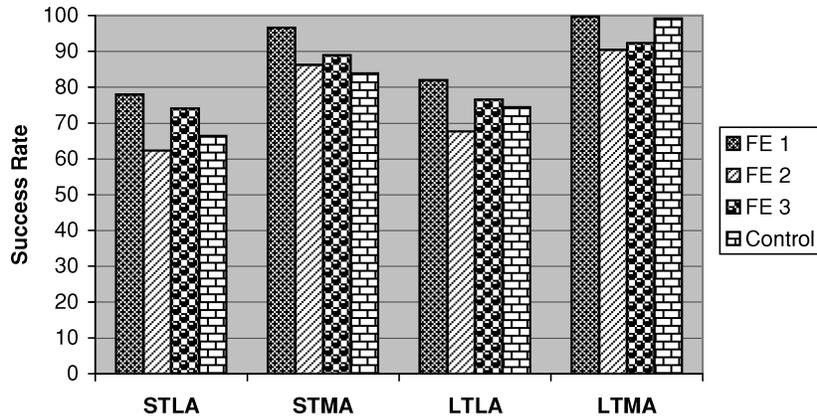


Fig. 8. Agent's performance in terms of success rate.

strategy, the cell corresponding to it is blank. The utilization of the different fitness functions reflects the varying behavior that the agent can employ in a given situation. It can be observed that the agents that utilised Fitness Equation 1 (where delivery is of utmost important) did indeed score a higher percentage in terms of success rate than the agents that used the other two fitness functions (see Figure 8) and the control model C , for all environments. Agents that used Fitness Equation 2 achieved the highest utility in all the environments (see Figure 9), whereas agents that used Fitness Equation 3 strike a balance between a high success rate and a high payoff. These results are very much as expected (see Section 4.2).

In the STLA environment based on Fitness Equation 1, the dominant strategy that emerged is the combination of the remaining time and the desperateness tactics ($w_{rt} = 0.76, w_{de} = 0.24$). In this particular situation, the agent's initial bids in both tactics are high and the agent quickly reaches its p_r ($k_{rt} = 0.73, \beta_{rt} = 99.59, k_{de} = 0.84, \beta_{de} = 56.04$). This behavior is rational since an agent that is interested in delivering the item successfully in this context should bid aggressively from the beginning to maximize its chances of acquiring the item. When Fitness Equation 2 is used, the dominant strategy that emerged is

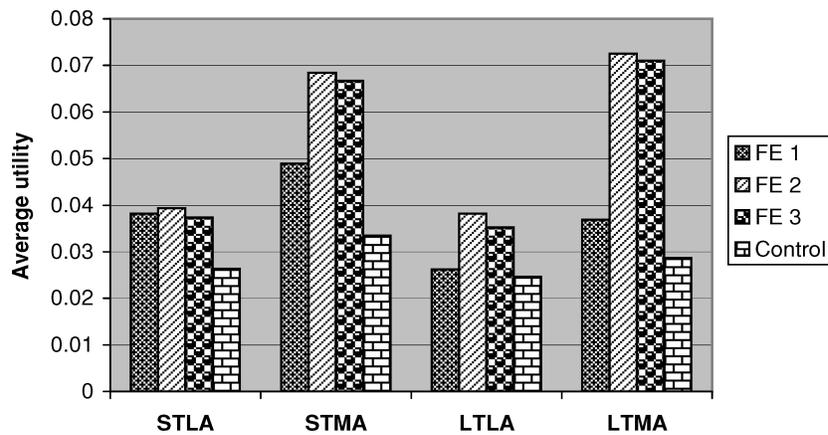


Fig. 9. Agent's performance in terms of average utility.

one that utilises all the tactics, but that places more importance on the remaining time and desperateness tactics. This is because the agent that is looking for a high payoff should consider the bargain tactic as one of the tactics to ensure a higher payoff. The strategy that emerged based on Fitness Equation 3 is one that considers the remaining auctions left and the desperateness tactics where the agent's initial bid is high and quickly reaches p_r . This strategy is similar to the one that emerged from Fitness Equation 1, but the rate at which it reaches p_r is slower. The reason for this is that an agent that is looking to maximize the payoff whilst ensuring delivery of the item needs to maintain a balance between a low bid price and the rate at which it reaches p_r .

In the STMA environment, an effective strategy should consider the remaining time and desperateness tactics highly since the allocated bidding time is limited (as per STLA). This is true when delivery of the item is important (as reflected in Fitness Equation 1's result), but also when payoff (refer to result of Fitness Equation 2) is the main consideration (here the agent combines all the tactics but heavier weights are placed on the desperateness and bargain tactics). This situation differs from STLA because here, the agent can afford to spend some time looking for a bargain since the number of active auctions is large. The dominant strategy that emerged based on Fitness Equation 3 is surprising because it combines the remaining time and the bargain tactics, instead of deploying a more aggressive behavior of combining the remaining time, desperateness and the bargain tactics. In this case, the strategy is aware of the large number of active auctions so it tries to get a higher payoff, but at the same time it takes into account the length of time it has left to bid.

The strategy that evolved for the LTLA environment based on Fitness Equation 1 is one that considers the remaining auctions and the desperateness tactics. This is because the strategy has to deliver the item successfully in an environment where there is a limited number of active auctions that the agent can participate in. As expected, when payoff is the main consideration, the strategy that evolved considers all tactics. The strategy that emerged

Table III. Strategies for STLA with Varying Private Valuations

p_r	(k_{rt}, β_{rt})	(k_{ra}, β_{ra})	(k_{ba}, β_{ba})	(k_{de}, β_{de})	$(w_{rt}, w_{ra}, w_{ba}, w_{de})$
68	(0.64, 5.44)	(0.57, 79.37)	(0.11, 0.15)	(0.75, 466.24)	(0.12, 0.00, 0.00, 0.88)
76	(0.05, 0.99)	(0.58, 508.59)	(0.23, 0.08)	(0.90, 86.53)	(0.00, 0.17, 0.00, 0.83)
82	(0.07, 97.11)	(0.06, 12.37)	(0.17, 0.46)	(0.75, 4.74)	(0.36, 0.35, 0.18, 0.11)

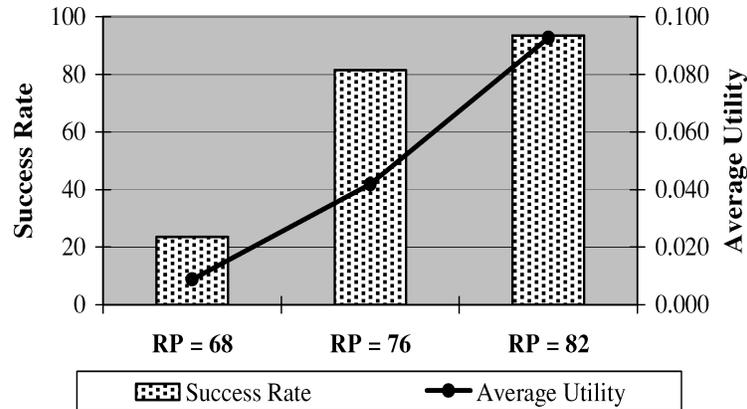


Fig. 10. Agent's performance with varying private valuations.

based on Fitness Equation 3 considers the remaining time, remaining auctions and desperateness tactics. Bargain is not considered here, since the number of active auctions in the marketplace is small (as per STLA).

All the strategies that evolved in the LTMA environment, for all fitness functions, achieved more than 90% success rate, but they differ in terms of payoff. The reason for this high success is due to the long bidding time, as well as the large number of active auctions that agents can participate in. Hence, the agent has many chances of winning. In this particular situation, the main consideration is the payoff. As can be seen, the strategies that utilised Fitness Equation 2 and 3 generate higher payoffs when compared to the strategy that evolved based on Fitness Equation 1 and the control model *C*. The reason for this is that both *C* and Fitness Equation 1 consider delivery as the most important criteria and payoff is not taken into account.

Turning now to the second class of experiments. Table III shows the strategies that evolved in the STLA environment based on Fitness Equation 3. Fitness Equation 3 is used here since it offers a reasonably high success and payoff. In this context, Figure 10 shows that the success rate and the payoff increase when p_r increases. The high payoff that the agent receives when using the strategy evolved with p_r indicates that the agent actively tries to look for bargain when private valuation is high ($p_r = 82$). In contrast, when the private valuation is low ($p_r = 68$), the agent evolves a strategy that combines the remaining time and the desperateness tactics to take advantage of the limited time, limited number of active auctions and limited p_r . The strategy that emerged with $p_r = 76$ is similar to the one that evolved with $p_r = 68$, but this time, it considers the remaining auctions left instead of the remaining time left. With a higher private

valuation, the agent has a better chance of obtaining the item, enabling it to switch to a strategy that focuses on the desperateness tactic and the remaining auctions left tactic. When p_r is high, the strategy that emerged considers all tactics as expected. These results led us to conclude that evolving the strategies for finer subdivisions (e.g., by further partitioning the STLA environment into more private valuation divisions) will better tune the agent's bidding strategy and result in a more superior performance.

In summary, several conclusions can be drawn from these results. First, GAs can be used to evolve strategies that are successful in particular environments (in all cases, performance is superior to the control model). Second, when selecting a strategy to bid in the multiple auctions environment, the agent needs to determine the current environment's type, as well as the user's preferences. Depending on these two, it can then decide which strategy to deploy. The result presented in Figure 10 shows that as the private valuation increases, the success rate increases, therefore allowing the strategy to deploy a bargaining behavior to generate a higher payoff. Finally, the results also indicate that the categories of environment for which the strategies need to be evolved can be further subdivided into finer divisions so that the agent can better tune its bidding strategy to its prevailing circumstances.

5. 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. This 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 environments into 54 categories (shown in Table IV) based on the key determinants identified in Section 4.4; namely the private valuation, the remaining time left, the remaining auctions left and the agent's behavior (whether looking for a bargain, desperate, or a mixture of both). The private valuation can be broadly categorised by its value: low (RP1), medium (RP2) or high (RP3). In this particular implementation, low is between 60 and 72, medium is between 73 and 79 and high is between 80 and 82. These three categories are then refined further into three subcategories based on the agent's behavior (FE 1 is desperate, FE 2 is looking for a bargain and FE 3 aims for a balance of both). Finally, each subcategorization is further divided into 54 grouping that consist of the short time less auctions environment (STLA), short time many auctions environment (STMA), medium time less auctions environment (MTLA), medium time many auctions environment (MTMA), long time less auctions environment (LTLA) and long time many auctions environment (LTMA). In this case, less auctions and many auctions are as defined in Section 4. Short time is defined as $10 \leq t_{\max} \leq 20$, medium time is defined as $21 \leq t_{\max} \leq 50$ and long time is defined as $51 \leq t_{\max} \leq 100$. As an example, RP1FE2STMA represents an environment with a low private valuation, where the bidding time is short, but where there are many auctions and the customer is interested in a bargain.

Table IV. The Environments

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
	LTMA	LTMA		LTMA				
	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

The categorization of the private valuation is made based on the auction closing price distribution (as per Section 3.3, the closing price mean is 76 and the standard deviation is 5). Fifty percent of the auctions should be won by bidders with medium private valuations, 25% by bidders with low private valuations and the remaining 25% by the bidders with high private valuation. 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 PriceSCAN,¹⁴ DealTime,¹⁵ and BottomDollar.¹⁶ 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 4.3 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

¹⁴<http://www.pricescan.com/>.

¹⁵<http://www1.dealtime.com/>.

¹⁶<http://www.bottomdollar.com/>.

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 (see Table V). Upon selection of the appropriate strategy the agent proceeds as defined in Figure 2.

5.1 Experimental Evaluation

The hypothesis that we seek to evaluate in this section is that our intelligent bidding strategy performs effectively in a wide range of bidding contexts. Here, the performance of the agent is measured in terms of success rate and average utility (as defined in Section 4.4). As our control models, we use an agent (*C1*) that has a single fixed strategy based on the user's behavior and an agent (*C2*) that possesses a random behavior. In particular, we picked the strategy that was evolved for the environments RP2FE1MTMA, RP2FE2MTMA and RP2FE3MTMA (as discussed in Section 4.4)¹⁷ for *C1*. Thus, *C1* has three different strategies; when a user is interested in a bargain, *C1* selects the strategy that was evolved for RP2FE2MTMA, when a user is desperate for the item, *C1* employs the 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. In particular, Table VI shows the values of the respective k , β , and the weights for each bidding constraint employed by *C1*. Control *C2*, on the other hand, is an agent that utilises a strategy in line with the user's preference that is randomly picked (for each run) from those listed in Table V. Thus, if a user is interested in a bargain, the agent randomly selects any of the 18 strategies in FE2. When a user is desperate for the item, the agent selects randomly from the strategies in FE1. Finally, when a user is looking for a combination of both, the agent selects the strategies in FE3 randomly. It was decided to use three strategies for the control models, rather than a single strategy, so that we could measure the performance of the agents in terms of the individual user's preferences. At any point in time, our intelligent agent will always possess one behavior (looking for a bargain, desperate or some combination of these two). Obviously, if its behavior is one that is interested in a bargain, the average utility should be high, but the success rate may be low. The converse of this is when our intelligent agent possesses desperate behavior; the success rate should go up, but the average utility may go down. Thus, to compare like with like, we consider the control model as similarly having the strategy that seeks to maximize the desired user behavior.

The experimental set up is broadly as described in Section 3.3. In particular, the agent and the control model were run 1000 times in the marketplace. ANOVA was used to test the hypothesis about the differences of the success rate means (when running the experiment 200, 400, 600, 800 and 1000 times) and the procedure revealed that for all experiments, the differences between means were not significant ($F_{4,15} = 0.134$, $p > 0.05$). Thus, the results obtained are statistically significant. The user's requirement was randomly allocated;

¹⁷We 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.

Table V. Strategies for the Intelligent Agent

Behavior	(k_{rt}, β_{rt})	(k_{ra}, β_{ra})	(k_{ba}, β_{ba})	(k_{de}, β_{de})	$(w_{rt}, w_{ra}, w_{ba}, w_{de})$
RP1FE1STLA	(0.88, 9.61)	(0.46, 60.63)	(0.14, 0.31)	(0.83, 315.43)	(0.03, 0.66, 0.00, 0.31)
RP1FE2STLA	(0.14, 21.82)	(0.02, 0.26)	(0.18, 0.17)	(0.75, 955.67)	(0.52, 0.00, 0.00, 0.48)
RP1FE3STLA	(0.24, 699.85)	(0.71, 6.86)	(0.22, 0.40)	(0.81, 7.72)	(0.04, 0.05, 0.00, 0.91)
RP1FE1STMA	(0.91, 290.24)	(0.93, 5.00)	(0.16, 0.46)	(0.83, 5.15)	(0.79, 0.00, 0.00, 0.21)
RP1FE2STMA	(0.64, 234.36)	(0.88, 0.99)	(0.22, 0.10)	(0.76, 6.06)	(0.54, 0.03, 0.03, 0.40)
RP1FE3STMA	(0.39, 97.17)	(0.82, 63.22)	(0.24, 0.34)	(0.76, 2.42)	(0.18, 0.51, 0.00, 0.31)
RP1FE1MTLA	(0.22, 64.16)	(0.29, 255.81)	(0.10, 0.21)	(0.78, 25.23)	(0.00, 0.77, 0.00, 0.23)
RP1FE2MTLA	(0.84, 975.47)	(0.92, 964.05)	(0.25, 0.45)	(0.84, 6.05)	(0.00, 0.06, 0.00, 0.94)
RP1FE3MTLA	(0.13, 527.00)	(0.92, 0.47)	(0.26, 0.31)	(0.78, 53.62)	(0.80, 0.18, 0.01, 0.01)
RP1FE1MTMA	(0.52, 325.34)	(0.52, 5.18)	(0.26, 0.24)	(0.83, 33.84)	(0.83, 0.00, 0.00, 0.17)
RP1FE2MTMA	(0.80, 9.06)	(0.25, 0.89)	(0.18, 0.23)	(0.88, 193.35)	(0.63, 0.00, 0.03, 0.34)
RP1FE3MTMA	(0.71, 300.22)	(0.83, 0.87)	(0.26, 0.15)	(0.86, 5.64)	(0.00, 0.01, 0.00, 0.99)
RP1FE1LTLA	(0.11, 61.68)	(0.67, 2.53)	(0.14, 0.42)	(0.72, 475.00)	(0.23, 0.00, 0.00, 0.77)
RP1FE2LTLA	(0.12, 604.92)	(0.77, 0.61)	(0.24, 0.11)	(0.85, 38.93)	(0.87, 0.13, 0.00, 0.00)
RP1FE3LTLA	(0.48, 945.65)	(0.75, 1.50)	(0.19, 0.38)	(0.85, 35.03)	(0.14, 0.14, 0.00, 0.72)
RP1FE1LTMA	(0.45, 31.57)	(0.97, 870.65)	(0.20, 0.37)	(0.73, 25.35)	(0.00, 0.44, 0.00, 0.56)
RP1FE2LTMA	(0.86, 0.19)	(0.26, 443.92)	(0.30, 0.47)	(0.79, 6.02)	(0.00, 0.06, 0.00, 0.94)
RP1FE3LTMA	(0.68, 2.64)	(0.43, 0.76)	(0.28, 0.38)	(0.87, 581.70)	(0.28, 0.00, 0.00, 0.72)
RP2FE1STLA	(0.63, 828.18)	(0.06, 586.30)	(0.17, 0.49)	(0.87, 83.56)	(0.69, 0.00, 0.00, 0.31)
RP2FE2STLA	(0.76, 0.60)	(0.38, 2.68)	(0.28, 0.13)	(0.73, 5.83)	(0.00, 0.00, 0.03, 0.97)
RP2FE3STLA	(0.94, 852.13)	(0.99, 9.99)	(0.24, 0.14)	(0.77, 24.34)	(0.00, 0.22, 0.05, 0.73)
RP2FE1STMA	(0.99, 0.83)	(0.92, 64.17)	(0.21, 0.24)	(0.84, 7.79)	(0.23, 0.77, 0.00, 0.00)
RP2FE2STMA	(0.16, 154.88)	(0.87, 0.92)	(0.17, 0.33)	(0.82, 8.56)	(0.63, 0.08, 0.09, 0.20)
RP2FE3STMA	(0.06, 1.85)	(0.93, 77.07)	(0.13, 0.11)	(0.76, 725.77)	(0.18, 0.39, 0.00, 0.43)
RP2FE1MTLA	(0.88, 535.04)	(0.27, 7.51)	(0.28, 0.22)	(0.75, 94.99)	(0.27, 0.00, 0.00, 0.73)
RP2FE2MTLA	(0.00, 9.38)	(0.94, 2.61)	(0.23, 0.35)	(0.85, 56.62)	(0.28, 0.72, 0.00, 0.00)
RP2FE3MTLA	(0.73, 8.33)	(0.74, 6.95)	(0.24, 0.38)	(0.82, 6.59)	(0.40, 0.19, 0.00, 0.41)
RP2FE1MTMA	(0.89, 32.56)	(1.00, 0.09)	(0.14, 0.21)	(0.82, 357)	(0.95, 0.05, 0.00, 0.00)
RP2FE2MTMA	(0.83, 2.41)	(0.83, 78.85)	(0.17, 0.12)	(0.76, 648.30)	(0.21, 0.60, 0.07, 0.12)
RP2FE3MTMA	(0.92, 963.71)	(0.16, 0.48)	(0.15, 0.13)	(0.84, 57.56)	(0.00, 0.08, 0.00, 0.92)
RP2FE1LTLA	(0.53, 94.45)	(0.46, 16.32)	(0.22, 0.38)	(0.79, 168.42)	(0.44, 0.00, 0.00, 0.56)
RP2FE2LTLA	(0.81, 81.56)	(0.5, 8.85)	(0.15, 0.37)	(0.84, 6.87)	(0.22, 0.00, 0.05, 0.73)
RP2FE3LTLA	(0.28, 0.53)	(0.62, 971.46)	(0.17, 0.41)	(0.89, 266.92)	(0.04, 0.02, 0.02, 0.92)
RP2FE1LTMA	(0.21, 596.56)	(0.67, 21.62)	(0.28, 0.16)	(0.73, 96.00)	(0.13, 0.05, 0.00, 0.82)
RP2FE2LTMA	(0.64, 328.86)	(0.54, 0.44)	(0.25, 0.14)	(0.78, 4.65)	(0.87, 0.11, 0.01, 0.01)
RP2FE3LTMA	(0.76, 4.41)	(0.71, 41.41)	(0.17, 0.34)	(0.89, 93.09)	(0.09, 0.65, 0.09, 0.17)
RP3FE1STLA	(0.17, 64.30)	(0.10, 9.12)	(0.23, 0.17)	(0.84, 83.46)	(0.72, 0.01, 0.00, 0.27)
RP3FE2STLA	(0.49, 4.68)	(0.89, 0.66)	(0.18, 0.42)	(0.80, 5.89)	(0.46, 0.45, 0.00, 0.09)
RP3FE3STLA	(0.43, 73.87)	(0.84, 61.46)	(0.11, 0.11)	(0.80, 5.43)	(0.06, 0.15, 0.06, 0.73)
RP3FE1STMA	(0.73, 8.23)	(0.90, 79.72)	(0.21, 0.38)	(0.78, 36.97)	(0.38, 0.00, 0.00, 0.62)
RP3FE2STMA	(0.19, 0.35)	(0.73, 160.09)	(0.11, 0.28)	(0.79, 5.17)	(0.13, 0.08, 0.00, 0.79)
RP3FE3STMA	(0.86, 1.38)	(0.81, 4.49)	(0.12, 0.27)	(0.86, 3.16)	(0.15, 0.49, 0.00, 0.36)
RP3FE1MTLA	(0.83, 97.05)	(0.89, 71.76)	(0.16, 0.37)	(0.84, 349.48)	(0.74, 0.00, 0.00, 0.26)
RP3FE2MTLA	(0.62, 65.60)	(0.91, 33.69)	(0.19, 0.31)	(0.81, 671.41)	(0.77, 0.06, 0.09, 0.08)
RP3FE3MTLA	(0.18, 34.38)	(0.60, 803.84)	(0.22, 0.39)	(0.88, 4.60)	(0.39, 0.00, 0.08, 0.53)
RP3FE1MTMA	(0.10, 98.30)	(0.96, 68.06)	(0.20, 0.36)	(0.78, 97.29)	(0.94, 0.06, 0.00, 0.00)
RP3FE2MTMA	(0.47, 8.08)	(0.80, 0.53)	(0.29, 0.41)	(0.79, 3.63)	(0.28, 0.45, 0.01, 0.26)
RP3FE3MTMA	(0.37, 8.45)	(0.74, 999.85)	(0.27, 0.20)	(0.72, 3.22)	(0.71, 0.02, 0.00, 0.27)
RP3FE1LTLA	(0.28, 931.97)	(0.30, 7.16)	(0.22, 0.21)	(0.85, 35.95)	(0.46, 0.00, 0.00, 0.54)
RP3FE2LTLA	(0.66, 0.01)	(0.86, 0.81)	(0.13, 0.43)	(0.74, 624.95)	(0.09, 0.39, 0.00, 0.52)
RP3FE3LTLA	(0.04, 80.77)	(0.29, 1.71)	(0.18, 0.49)	(0.87, 927.76)	(0.96, 0.00, 0.00, 0.04)
RP3FE1LTMA	(0.21, 89.29)	(0.33, 69.06)	(0.17, 0.29)	(0.71, 55.04)	(0.16, 0.00, 0.00, 0.84)
RP3FE2LTMA	(0.30, 83.89)	(0.41, 13.55)	(0.30, 0.3)	(0.78, 33.51)	(0.46, 0.42, 0.12, 0.00)
RP3FE3LTMA	(0.31, 0.73)	(0.85, 2.97)	(0.14, 0.20)	(0.73, 2.89)	(0.03, 0.83, 0.00, 0.14)

Table VI. Strategies for the Control Model C1

Behavior	(k_{rt}, β_{rt})	(k_{ra}, β_{ra})	(k_{ba}, β_{ba})	(k_{de}, β_{de})	$(w_{rt}, w_{ra}, w_{ba}, w_{de})$
RP2FE1MTMA	(0.89, 32.56)	(1.00, 0.09)	(0.14, 0.21)	(0.82, 3.57)	(0.95, 0.05, 0.00, 0.00)
RP2FE2MTMA	(0.83, 2.41)	(0.83, 78.85)	(0.17, 0.12)	(0.76, 648.30)	(0.21, 0.60, 0.07, 0.12)
RP2FE3MTMA	(0.92, 963.71)	(0.16, 0.48)	(0.15, 0.13)	(0.84, 57.56)	(0.00, 0.08, 0.00, 0.92)

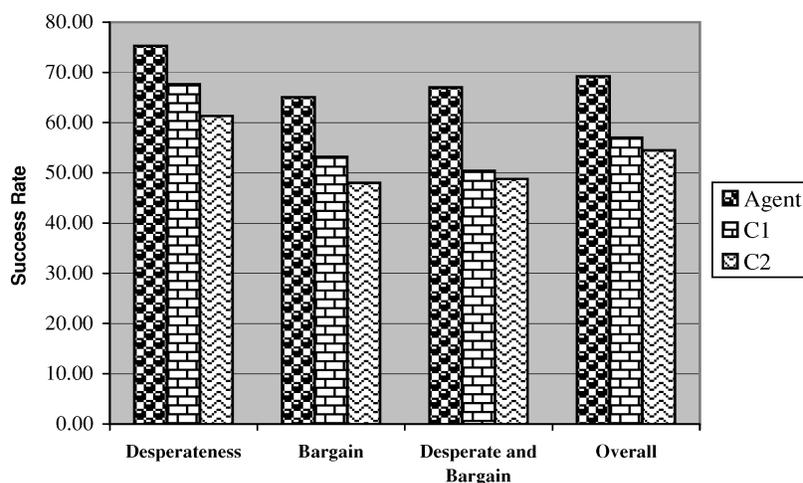


Fig. 11. Success rate comparisons.

the private valuation ranges from 70 to 82 and the time allocated ranges from 10 to 100.¹⁸ The user's intention is also generated randomly. The number of auctions running in the marketplace is between 2 and 60 and, as before, there are between 2 and 10 participants in each such auction.

The performance of the intelligent agent and the control model in terms of success rate is shown in Figure 11. The experiment is divided into four groups. The first three groups show the detailed performance of the agent and the control models based on a single behavior (desperateness, bargain, and both) and the last group shows the overall performance when all three behaviors are considered. It can be seen that our intelligent agent achieved a higher success rate for all the individual user behaviors and for the overall behavior. In desperateness mode, the intelligent agent achieved a 7% higher success rate than C1 and a 14% higher success rate than C2. This shows that by having the ability to change the strategy in accordance with the user's preference and the environment it is situated in, our agent can maximize its chances of succeeding. This is different from C1, in which a fixed strategy is used and where the agent views the environment and user's preferences as static. In this case, C1 was only successful in the environment for which it was evolved (RP2FE1MTMA). C2 achieves a lower success rate when compared to our intelligent agent and C1. Since the strategy selection in C2 is random, the strategy selected may not be suitable for the environment that the agent finds itself in. The situation

¹⁸This is slightly different from the two previous experiments in that the private valuation and the time are now generated randomly (instead of fixing them to a particular value).

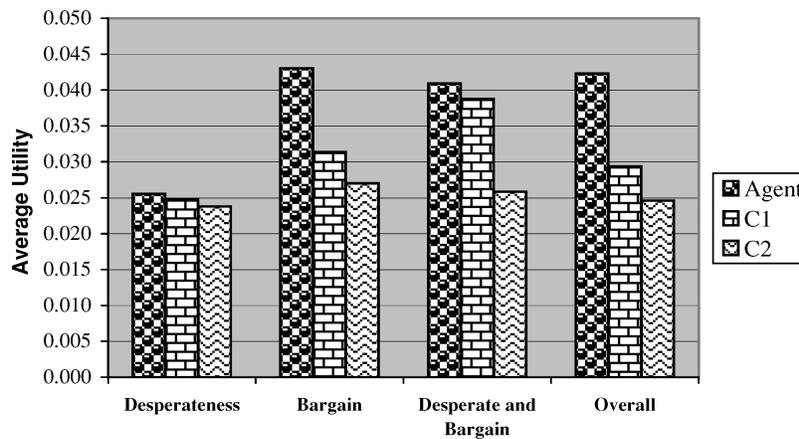


Fig. 12. Average utility comparisons.

is similar when our intelligent agent is in bargain mode. It achieved a higher success rate when compared to the two control models. When both bargain and desperation are considered, our intelligent agent also achieved a higher success rate than the control models. In both cases, *C1*'s performance is lower because its fixed strategy is inappropriate. *C1* starts bidding at a low value and slowly reaches its private valuation. Such a strategy is simply not effective in environments where there is a short bidding time, low private valuations or few active auctions in the marketplace. *C2* picks its strategy without taking into account the nature of its environment, resulting in a lower performance when compared to both *C1* and our intelligent agent. When all these behaviors are combined, our intelligent agent achieved 12% higher success rate than the control models.

Figure 12 shows the average utility obtained by the intelligent bidding agent and the control models. It can be seen that our agent performed well compared to the control model and it achieved a high average utility in all cases. However, in the desperateness mode, our agent's average utility surpassed the control models' average utility by only a small value (0.00077 more than *C1* and 0.00175 more than *C2*). The reason for this can be attributed to the strategies used in *C1* and *C2*. The former uses a fixed strategy that enables it to perform well in the RP2FE1MTMA environment and possibly the RP3 environments. This is because with a higher private valuation, *C1* has more chances of obtaining the item irrespective of its strategy. In this case, when it wins the auction, it will definitely obtain a high payoff resulting in a high average utility. In the latter case, when the private valuation falls into the RP3 category, the strategy picked yields a higher payoff than the others and this also results in a higher average utility. This is not the case for our intelligent agent because it uses a strategy that is evolved based on the current environment where it is not interested in obtaining a higher utility (rather it is interested in the delivery of the required item).

In summary, these results confirm our motivation's hypothesis that our intelligent bidding strategy can perform effectively in a wide range of environments.

It achieved a significantly higher success rate and a higher average utility in all the cases considered.

6. RELATED WORK

There have been several attempts to design sophisticated and efficient bidding strategies for agents participating in online auctions. For example, Faratin et al. [1998] is broadly similar to the mechanism defined in this article. 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, Sierra, and Jennings [2001] who analyzed the evolution of the negotiation strategies using GAs, and determined which of them are appropriate in which situations. The aim of this work was to perform an evaluation of the range of negotiation strategies by analyzing their relative success, and how these strategies evolve over time to become a fitter population. This approach is somewhat similar to our work, but the main difference is in the domain that we are dealing with (multiple auctions versus bilateral negotiations).

BiddingBot is a multi-agent system that supports users in attending, monitoring and bidding in multiple auctions through a process called co-operative bidding [Ito et al. 2000]. This approach demonstrates how agents can cooperate and work together to do the bidding process in multiple auctions. It consists of one leader and several bidder agents, where the leader agent acts as the coordinator and the facilitator of the whole bidding process. Bidding is done by exchanging messages between the user, the leader agent and the bidder agents. However, the main problem with this approach is that the agents do not actually make the bidding decision. This decision is left to the user. Thus, the agents do not have full autonomy and the decision-making process is slow since the agent needs to interact with the user from time to time.

The trading agent competition (TAC)¹⁹ provided a platform for agent designers to develop autonomous agents that can compete with one another in multiple simultaneous auctions for complimentary and substitutable goods. The key feature of TAC is that it required autonomous bidding agents to buy and sell multiple interrelated goods in auctions of different types [Greenwald and Stone 2001]. Each participating agent is a simulated travel agent with the goal of assembling a number of travel packages for its eight clients. Each client is characterised by a random set of preferences for the possible arrival and departure dates, hotel rooms and entertainment tickets. The objective of a TAC agent is to maximise the total satisfaction of its customers (i.e., the sum of the customer's utilities). The competition attracted a number of alternative agent designs (e.g., ATTac-2000 [Stone et al. 2001], RoxyBot [Boyan and Greenwald 2001], Aster [Greenwald and Stone 2001] and SouthamptonTAC [He and Jennings 2003]). Although there are clearly some similarities with our scenario, there are also a number of important differences. In particular, we concentrate on the bidding strategies to obtain a single item rather than worrying

¹⁹<http://www.sics.se/tac/>.

about the complementary goods that need to be bundled with the desired item. Moreover, in TAC, the type of auctions that are conducted are continuous one-sided auctions, standard English ascending multi-unit auctions and continuous double auctions as opposed to our environment that runs simultaneous standard English auctions, Dutch auctions and Vickrey auctions.

Preist et al. [2001] 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 [2001] also considers this environment, but utilizes 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 the multiple English auctions and their algorithms are not applicable in a heterogeneous protocol context. Byde et al. [2001] presented another decision theoretic framework that an autonomous agent can use to bid effectively across multiple auctions with various protocols (namely, English, Dutch, first price sealed bid and Vickrey auctions). In order to come up with the best bid value that guarantees the delivery of the item, an agent must always speculate about future events. To do this, Byde presented an approximation function that provides an estimate of the expected utility of participating in the set of future auctions. The decision making algorithm works in this way; it selects all the Dutch, English, and sealed bid auctions that the agent wishes to consider. It then tests the union of all three sets of bids to determine the utility of this course of action (using the approximation function). This process is repeated for all the possible combinations of Dutch, English and sealed bid auctions and returns the auction set with the highest expected utility. This auction set contains the list of all the auctions that the agent should bid in at that particular point in time. This approach can be employed to purchase single or multiple items in online auctions. 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 methods.

7. CONCLUSIONS AND FUTURE WORK

This article presented a novel bidding algorithm that can be used for an agent to participate in multiple online auctions, with multiple protocols, and varying start and end times. The bidding algorithm itself is based on multiple tactics, that each deal with a single facet of the agent's reasoning. These tactics are combined in order to give the agent's overall view, at a given moment in time, based on a systematic empirical evaluation. In the evaluation, it was discovered that the performance of the agent's bidding strategy is mainly dependent on its private valuation, the number of auctions left and the remaining time left. Hence, the bidding context can be categorized according to the user's preferences and the environment. Given this finding, we used genetic algorithms to evolve effective bidding strategies for particular types of preference in particular,

classes of environment. These evolved strategies were then embedded into the agent's reasoning strategy so that it can perform successfully across a wide range of auctions and can better tune its bidding behavior to its prevailing circumstances.

The model and method that we have used in this work can be readily extended or applied to different auction settings. For instance, if there are other bidding constraints that should affect the behavior of the agent, then these can easily be incorporated as a new family of tactics. Similarly, we can reclassify the environment to reflect any changes in the marketplace setting. As an example, in this work, the private valuation is broadly categorized by its value (low, medium or high). This can easily be generalized to cater for different type of items with varying prices, simply by mapping the mean price of a given item with the observed closing price distribution and regenerating the price ranges for the low, medium and high private valuations. This generalization can also be applied when subdividing the environment based on the number of auctions in the marketplace and the remaining time left. If there is a drastic change in the environment, these classifications can be expanded or shrunk to suit the current situation.

In terms of the practical application of our method, one of the key problems that needs to be addressed is how well an agent can assess the type of environment it is situated in. In the experiments reported in this article, we assume that the agent is able to quickly and accurately determine its environment type. However, we believe this will be significantly more difficult in real world scenarios. This, in turn, may require the agent to select a strategy probabilistically according to its beliefs at a given moment in time and then to alter this strategy as its beliefs change. This is an area that requires further investigation.

A second aspect that requires further consideration is the selection of the target bid from the list of potential auctions (Section 3.2). As it currently stands, the auction with the highest expected utility is selected as the agent's target auction. However, the expected utility is simply the product of the probability of the agent winning in that auction at a given bid value and the value of the agent's utility function. This probability is calculated based on the closing price distribution for all the auctions observed over a given time. In this setting, this information is essentially static. Moreover, we are dealing with a group of bidders that have broadly similar characteristics (they use dominant strategies to their respective one-shot auctions and their private valuations are drawn from the same probability distribution). In real market settings, the behaviors of the bidders are likely to be more complex (e.g., they may employ variable bidding strategies and their private valuations may be based on their own private assessment of the item or the private valuations of their competitors). In such cases, more accurate probability prediction is necessary to ensure that the right target auction is picked. This should also include updating the bid histories from time to time (the agent should also consider auctions that are closing while it is still active in the marketplace), as well as taking into account the behaviors of the other auction participants (or what the agent believes about its competitors' behaviors). Some work has been undertaken in this area (such as belief-based learning [Preist et al. 2001] and probability estimation method

[Dumas et al. 2002]), but the probability estimations are calculated based on English and sealed bid auctions only.

There are also several other areas that also require further investigation. First, given the success of our bidding strategy, we believe that it may be widely adopted. Thus, we need to determine what happens to individual and system performance when there are multiple agents using our strategy in a given auction context. Second, this algorithm is designed to purchase a single item from the marketplace. We would like to determine whether the algorithm can be extended to deal with the more general case where the agent needs to purchase multiple items of a given type. Finally, we would like to compare our heuristic method with the more principled decision theoretic approach of Bye et al. [2002] to determine the relative strength and weaknesses of the two methods.

ACKNOWLEDGMENTS

We wish to thank Chris Preist for his comments and suggestions and Viet Dung Dang for his help in formalizing some of the decision making procedures in the early stage of this work.

REFERENCES

- ANTHONY, P., DANG, V. D., HALL, W., AND JENNINGS, N. R. 2001. Autonomous agents for participating in multiple online auctions. In *Proceedings of the IJCAI Workshop on E-Business and Intelligent Web*. 54–64.
- ANTHONY, P. AND JENNINGS, N. R. 2002. Evolving bidding strategies for multiple auctions. In *Proceedings of the 15th European Conference on Artificial Intelligence*. 178–182.
- BAPNA, R., GOES, P., AND GUPTA, A. 2001. Insights and analyses of online auctions. *Commun. ACM* 44, 11 (Nov.), 43–50.
- BEASLEY, D., BULL, D. R., AND MARTIN, R. R. 1993a. An overview of genetic algorithms: Part I, Fundamentals. Tech. Rep. 15(2), University of Purdue.
- BEASLEY, D., BULL, D. R., AND MARTIN, R. R. 1993b. An overview of genetic algorithms: Part II, Research topics. Tech. Rep. 15(4), University of Purdue.
- BLICKLE, T. AND THIELE, L. 1995a. A comparison of selection schemes used in genetic algorithms. Tech. Rep. 11, TIK.
- BLICKLE, T. AND THIELE, L. 1995b. A mathematical analysis of tournament selection. In *Proceedings of the 6th International Conference on Genetic Algorithm (ICGA95)*. Morgan-Kaufmann, San Francisco, Calif., 9–16.
- BOYAN, J. AND GREENWALD, A. 2001. Bid determination in simultaneous actions: An agent architecture. In *Proceedings of the 3rd ACM conference on Electronic Commerce*. ACM, New York, 210–212.
- BYE, A. 2001. Programming model for algorithm design in simultaneous auctions. In *Proceedings of the 2nd International Workshop on Electronic Commerce*. ACM, New York, 152–163.
- BYE, A., PREIST, C., AND JENNINGS, N. R. 2002. Decision procedures for multiple auctions. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and MultiAgent Systems: Part 2*. ACM, New York, 613–620.
- COHEN, P. R. 1995. *Empirical Methods for Artificial Intelligence*. The MIT Press, Cambridge, Mass.
- DUMAS, M., GOVERNATORI, G., HOFSTEDE, A. T., AND RUSSELL, N. 2002. An architecture for assembling agents that participate in alternative heterogeneous auctions. In *Proceedings of the Workshop on Research Issues in Data Engineering*. IEEE Computer Society Press, Los Alamitos, Calif., 75–83.
- FARATIN, P., SIERRA, C., AND JENNINGS, N. R. 1998. Negotiation decision functions for autonomous agents. *Int. J. Robot. Autonom. Syst.* 24, 3-4, 159–182.

- GREENWALD, A. AND STONE, P. 2001. Autonomous bidding agents in the trading agent competition. *IEEE Internet Comput.* 52–60.
- HE, M. AND JENNINGS, N. R. 2003. SouthamptonTAC: An adaptive autonomous trading agent. *ACM Trans. Internet Tech.* 3, 3 (Aug.), 000–000.
- HE, M., JENNINGS, N. R., AND LEUNG, H. F. 2003. On agent-mediated electronic commerce. *IEEE Trans. Knowl. Data Eng.* 15, 4 (July/Aug.).
- ITO, T., FUKUTA, N., SHINTANI, T., AND SYCARA, K. 2000. Biddingbot: A multiagent support system for cooperative bidding in multiple auctions. In *Proceedings of the 4th International Conference on MultiAgent Systems* (Boston, Mass.). IEEE Computer Society Press, Los Alamitos, Calif., 399–400.
- KLEMPERER, P. 1999. Auction theory: A guide to literature. *J. Econ. Surv.* 13, 3, 227–286.
- MATOS, N., SIERRA, C., AND JENNINGS, N. R. 2001. Successful negotiation strategies: An evolutionary approach. In *Proceedings of the 3rd International Conference on Multi Agent Systems* (Paris, France). IEEE Computer Society Press, Los Alamitos, Calif., 182–189.
- MCAFEE, R. P. AND McMILLAN, J. 1987. Auctions and bidding. *J. Econ. Lit.* 25, 639–738.
- MCAFEE, R. P. AND McMILLAN, J. 1992. Bidding rings. *Amer. Econ. Rev.* 82, 3 (June), 579–599.
- MITCHELL, M. 1996. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, Mass.
- PREIST, C., BARTOLINI, C., AND PHILLIPS, I. 2001. Algorithm design for agents which participate in multiple simultaneous auctions. In *Agent Mediated Electronic Commerce III*. Lecture Notes in Artificial Intelligence. Springer, Berlin, Germany, 139–154.
- SANDHOLM, T. W. 1999. Distributed rational decision making. In *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, G. Weiss, Ed. The MIT Press, Cambridge, Mass., USA, 201–258.
- STONE, P., LITTMAN, M. L., SINGH, S., AND KEARNS, M. 2001. Attac-2000: An adaptive autonomous bidding agent. In *Proceedings of the 5th International Conference on Autonomous Agents*. ACM, New York, 238–248.
- VICKREY, W. 1961. Counterspeculation, auctions and competitive sealed tenders. *J. Fin.* 16, 8–37.

Received October 2002; revised February 2003; accepted February 2003