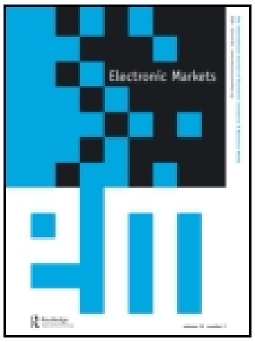

Electronic Markets

ISSN: 1019-6781 (Print) 1422-8890 (Online) Journal homepage: http://www.tandfonline.com/loi/rema20

Designing Pricing Mechanisms for Autonomous Agents Based on Bid-Forecasting

Dionisis D. Kehagias , Andreas L. Symeonidis & Pericles A. Mitkas

A b s t r a c t

Autonomous agents that participate in the electronic market environment introduce an advanced paradigm for realizing automated deliberations over offered prices of auctioned goods. These agents represent humans and their assets, therefore it is critical for them not only to act rationally but also efficiently. By enabling agents to deploy bidding strategies and to compete with each other in a marketplace, a valuable amount of historical data is produced. An optimal dynamic forecasting of the maximum offered bid would enable more gainful behaviours by agents. In this respect, this paper presents a methodology that takes advantage of price offers generated in e-auctions, in order to provide an adequate short-term forecasting schema based on time-series analysis. The forecast is incorporated into the reasoning mechanism of a group of autonomous e-auction agents to improve their bidding behaviour. In order to test the improvement introduced by the proposed method, we set up a test-bed, on which a slightly variant version of the first-price ascending auction is simulated with many buyers and one seller, trading with each other over one item. The results of the proposed methodology are discussed and many possible extensions and improvements are advocated to ensure wide acceptance of the bid-forecasting reasoning mechanism.

Keywords: e-marketplaces, auctions, autonomous agents, bidding, forecasting

Designing Pricing Mechanisms for Autonomous Agents Based on Bid-Forecasting

DIONISIS D. KEHAGIAS, ANDREAS L. SYMEONIDIS AND PERICLES A. MITKAS

A u t h o r s

Dionisis Kehagias
(diok@ee.auth.gr) is a PhD student affiliated with the Department of Electrical and Computer Engineering of the Aristotle University of Thessaloniki, and also a Research Associate at the Informatics and Telematics Institute of the Centre of Research and Technology Hellas (ITI/CERTH), in Greece. His research interests include autonomous agents, multiagent systems, agent mediated e-commerce and automated negotiations.

Andreas L. Symeonidis
(asymeon@ee.auth.gr) is a research associate at the Department of Electrical and Computer Engineering of the Aristotle University of Thessaloniki. He holds a PhD degree in Computer and Electrical Engineering from the same department in 2004. He is also affiliated with ITI/CERTH. His current research interests include data mining techniques for improving agent intelligence and machine learning.

Pericles A. Mitkas
(mitkas@eng.auth.gr) is an Associate Professor with the Department of Electrical and Computer engineering at the Aristotle University of Thessaloniki and a faculty affiliate of ITI/CERTH, Greece. His research interests include parallel architectures for large multimedia databases, data mining and data warehousing, intelligent agents, e-commerce and bioinformatics.

INTRODUCTION

Software agents that participate in electronic business transactions in an effort to increase revenue for humans introduce a new form of automatic negotiations. These agents enjoy a large degree of autonomy because, despite being computer programs, they actually undertake the responsibilities of rational reasoning and deal making on behalf of humans. In this way, agents can mitigate the difficult task of having to deliberate over offering the best price for purchasing a particular item. The successful performance of agents in trading scenarios, typically reflected on the amount of money they earn or save, ensures that agent-based electronic marketplaces can adequately replace the ordinary ones. Even though humans seem to negotiate using complex procedures (Beam and Segev 1997), it is not feasible for them to monitor and understand the attributes of a sequence of negotiations in a large competitive environment, such as the one that hosts electronic auctions (e-auctions). Besides their superiority in terms of monitoring and 'remembering', agents are also able to follow a specific course of action towards their effort to increase their profit

Volume 15 (1): 53–62. www.electronicmarkets.org DOI: 10.1080/10196780500035340

without being diverted by emotional influence as humans do. Inspired by these advantages of agents in electronic marketplaces, we deployed a new technique for the exploitation of available data, in order to improve the bidding mechanism of agents for e-auctions.

An interesting issue concerning agents in e-commerce is the creation of both rational and efficient agent behaviours, to enable reliable agent-mediated transactions. In our research work we focus on the improvement of agent behaviours in auctions. This achievement becomes quite appealing as it can significantly advance the quality of existing agent-mediated e-commerce applications. The problem of agents in complex auctions with many participants is the lack of information possessed by their rivals (Bichler 2000). Such information is generated only in the form of bids publicly revealed in the auction. Whenever an agent announces its bid, it passes some information on its bidding behaviour to its competitors. But this information cannot reveal the details of its private bidding mechanism, or its valuation of the item being auctioned. However, if an external observer watches closely all bids publicly submitted over time, he may identify a tendency in the bid curve, which is unique for each particular auction. Appropriate analysis of the data produced as an auction progresses (historical data) can lead to short-term forecasting of the next bid with sufficient accuracy. This information may then be exploited by agents to create successful bidding behaviours. For example, agents may bid close to the forecast value. The core set of tools used for forecasting are provided by time-series analysis. In particular, the autoregressive model (Hamilton 1994) of order two, also denoted as AR(2), is used for estimating the next highest bid. The main objective of our research work is to prove that the results of the proposed data analysis can benefit the bidder's outcome. Indeed, as it is shown experimentally later on in the paper, given that the prediction error is low, the proposed bidding mechanism yields higher profits for buyers.

In order to evaluate the proposed method, we created an auction simulation environment, on which we can apply the short-term bid forecasting mechanism on a sequence of English-like first-price ascending auctions that close at different times, in an iterative manner. A single buyer, who adopts the forecasting schema, is benchmarked in the simulation environment against a set of other agents who employ a simple bidding mechanism, without taking into account the generated forecasting. The experiments show that the inclusion of forecasting in the bidding mechanism of agents can ensure improvement of their bidding behaviour.

RELATED WORK

In an auction environment with many agents a large amount of data are produced as bidders submit their offers. While the history of an auction may contain significant information, it was never exploited for short-term forecasting of the next bid by any of the known bidding techniques. On the contrary, a common set of methodologies, such as the ones described in (Greenwald and Stone 2001), use data from a series of the same or similar auctions closed in the past, in order to perform long-term forecasting of the winning bid. There is a key difference between the two approaches. In the former case, bid forecasting is based on data generated within the same auction, while in the latter the forecasting is based on values corresponding to the winning bid (the one that remains after the closing of a series of past auctions).

Different approaches have been proposed in the recent literature, which provide efficient bidding mechanisms for autonomous agents. Work that is more similar to the one presented in this paper is the development of techniques for price prediction (Wellman *et al.* 2003b). In the recently held agent trading competition (TAC) (Wellman *et al.* 2003a), agents deploy different bidding strategies in order to achieve the best price for travelling services, such as hotel accommodation, in order to satisfy the preferences of a group of clients.

The sizeable research effort on the development of soft computing models for improving trade-offs in negotiations between agents was a major source of inspiration for our work (Faratin *et al.* 2000). Even though these methodologies are more suitable for bilateral negotiations involving many negotiating attributes, which is a problem of higher complexity compared to auctions of a single item, they rely on the same mechanism for calculating the next offer. Other similar techniques have been proposed (He *et al.* 2002) and applied on different types of auctions, including continuous double auctions. These auctions involve multiple agents that look for services provided by multiple agents at any one time. A detailed survey of the state of the art in agent-mediated e-commerce can be found in (He *et al.* 2003).

THE AUCTION ENVIRONMENT

Prior to presenting the bidding forecasting methodology, it is important to define the environment that the agents perceive and act upon. In our implementation we deal with the situation of time-bound auctions with multiple buyers, $b_1, \ldots b_N$, a single seller s and one item g to be auctioned. This configuration is shown in Figure 1. The auction simulation environment is implemented as a multiagent system where agents participate in bilateral transactions, such as the ones described in (Kurbel and Loutchko 2002), even though the latter adopt different bidding techniques. This section describes the protocol that determines the rules of the auctions, as well as the details of the bidding rules of buyers.

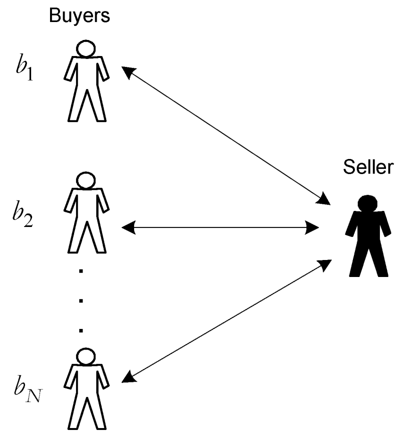

Figure 1. The interaction model among auction participants

Auction protocol

Two naming conventions are followed throughout this paper to describe the two distinct classes of participating agents: the terms '*buyers*' and '*sellers*' are used to determine all agents whose goal is to buy or sell goods in the auction, respectively. Other agents playing facilitating roles, which may exist in the multiagent marketplace, are ignored in this context. In the environment under consideration we assume monopoly from the seller's side, while competition between potential multiple sellers of the same item (McAffee and McMillan 1987) is not discussed in this paper. The price, that a buyer offers in order to acquire the desired item, is called a *bid*. Let β be the highest current bid in the auction, which will be referred to as *outstanding bid*. Given that item g is unique, at any particular time of the auction there is only one outstanding bid for g.

To demonstrate the bid forecasting mechanism, we have used the time-bound first-price ascending case. This choice was made because short-term forecasting needs historical data generated within the auction. In other words, the auction has to reveal intermediate information as it progresses. Auction types with this property include the English first-price outcry auction and its permutations. On the other hand, short-term forecasting cannot be applied on auctions that do not reveal intermediate bid information, commonly known as 'sealed' auctions. Examples of sealed auctions include the Dutch (Sandholm 1999), and Vickrey (Vickrey 1961) auctions. Additional constraints imposed in the auction environment are described as follows: First, the auction has a limited duration t_{max}, after which the auction clears, and the agent with the highest bid purchases the auctioned item and wins the auction. This constraint is met in several real online auction environments such as eBay (http://www.ebay.com). Second, a minimum acceptable increment step $\Delta\beta$ is imposed to all bidders. Third, all agents submit their bids but only the current highest bid (price quote) is announced to them.

In the context of the simulated auction, we assume that the competitive agents are risk-averse. This ensures that no agent will act irrationally in order to purposefully win the auction, even damaging its profit. For example, a buyer who is risk-averse does not overbid too early in the auction, and keeps bidding at 'reasonable' prices. The implication of this constraint in common-value auction models is straightforward. Another reason that prevents rational agents from overbidding is the awareness of the 'winner's curse' (Milgrom 1989), which enforces a winner who overbids to pay more for an item than its true value.

We also assume that each buyer b_i has a valuation function υ_i that maps the value of the offered price x at a given time t to the interval $[0,1]$. The valuation can be modelled as a two-variable function $\upsilon_i\ (x,\ t)$. The argument t reflects that the valuation also depends on the time remaining until the end of the auction. For instance, a buyer may have abstained for a long period since the beginning of the auction but he starts submitting new bids as the auction approaches its deadline (conservative behaviour).

We use three definitions to formally describe the auction mechanism.

Definition 1: Let $B=\{b_1,\ b_2,\ \ldots,\ b_N\}$ be the set of buyer identifiers, s the instance of the seller, $\boldsymbol{\upsilon}_B=\{\upsilon_1(\chi,\ \tau),\ \upsilon_2(\chi,\ \tau),\ \ldots,\ \upsilon_N(\chi,\ \tau)\}$ the vector of valuation functions for each buyer (buyer b_i has valuation υ_i, $1 \leqslant i \leqslant N$) with χ the current offered price for item g at time τ, and let t_{max} be the duration of the auction. An *instance* of the auction at time t can be described by the tuple:

$$a=\langle B,\ \upsilon_B,\ s,\ \beta,\ g,\ t,\ t_{max}\rangle$$

In this definition time t has a discrete sense. That is, even though the world evolves in a continuous manner, we let the agents perceive it in discrete-time steps. The discrete time is used for practical reasons, because the methodology proposed in this paper requires many tuples of time-series data, which are to be sampled in equal time intervals.

Definition 2: A *transaction* between a buyer b and a seller s, is the process described by the pseudo-code shown in Figure 2.

According to the function transaction, we let seller s inform any interested buyer about the current highest bid. Each buyer can then either submit a new bid x, calculated according to its private bidding tactic, or wait until the offered price reaches his own valuation. On bid submission, the *transaction* function returns the Boolean constant *true*, otherwise it returns *false*.

Definition 3: An *auction*, organized by a single seller s, is the process defined in Figure 3 in pseudo-code.

In the procedure auction, the seller s waits for a new message *msg* to arrive from the buyers' side until the

```
function transaction (Buyer b, Seller s) BEGIN
  static: β, t         /* β: outstanding bid, t: current time */
  deal ← false;
  β ← s.get_outstanding_bid;
  x ← b.calculate_new_bid_according_to_tactic;
  if (b.strategy_allows) then
    b.submit_a_new_bid;
    deal ← true;
  endif
  return deal;
END
```

Figure 2. Algorithm for a transaction between a buyer and a seller

auction terminates ($t \geqslant t_{max}$). As soon as *s* receives a valid new message (*msg*≠NULL), a communication activity occurs between this buyer and the seller. From this point on, there are two possibilities: a) either the buyer *b* submits a new bid, or b) decides to wait until the current outstanding bid meets his own valuation. In the former case, *b* is considered a potential winner of the auction until a new higher bid is submitted, or until the auction closes. Otherwise, if the buyer submits no new bid, nothing happens until the seller receives a new message. This process iterates until the auction terminates and the buyer who last offered the highest bid is declared as the winner of the auction. Then, the seller informs all participants that the auction has terminated and announces the final winner.

Following the auction protocol described by definitions 1 to 3, an auction can begin, proceed, and terminate in the e-marketplace environment. These three phases that determine the lifecycle of an auction are summarized in the following steps:

- *Step1*. The auction initiates. The seller announces the item *g* to be auctioned. The initial outstanding bid is set to seller's initial desired price and time is reset: $t \leftarrow 0$.
- *Step2*. The auction proceeds. Procedure *auction* introduced in definition 3 is executed.
- *Step3*. The auction terminates. Seller announces the winner of the auction and procedure *auction* returns.

Having defined the auction attributes, we can get into the details of the decision mechanisms deployed by the involved agents.

Valuation functions

The role of the valuation function is to preserve the rational behaviour of a buyer. We assume that buyers only take into account two issues in order to calculate their valuation of the offered item: the *price* of the auctioned item and the *time* remaining until the end of the auction. This behaviour is controlled by a – unique to each agent – valuation function that maps the value of the qualitative attributes of the item to be purchased to a real number in the interval [0, 1]. After a buyer has evaluated the value of the offered price, he can decide to bid or not.

The right choice of the valuation function is a matter of good auction design. For a buyer b_i to decide on whether to submit a new bid or not, he must evaluate the offered bid according to the following type of

```
procedure auction BEGIN
  static: t, s                            /* t: current time, s: seller */
  constant: tmax                             /* deadline of the auction */
  deal ← false;
  time_starts_counting;
  repeat               /* seller waits for receiving a new message */
    msg ←  s.receive_message_from_buyer;
    if (msg ≠ NULL) then                     /* if a new message comes up
   */
      b ← msg.get_sender;                 /* fetch the corresponding buyer */
      deal ← transaction (b, s);                     /* invoke transaction */
      if (deal = true) then                 /* if buyer b submits a new bid */
       winner ← b;                          /* b becomes a potential winner */
      endif
    endif
  until t ≥ tmax              /* repeat until auction reaches its deadline */
  seller_announces_the_winner;
END
```

Figure 3. The main auction procedure

scoring functions taken from a bilateral negotiation model introduced in (Faratin *et al.* 2000):

$$\upsilon_i^p(x) = \left(\frac{x_i^{\max} - x}{x_i^{\max} - x_i^{\min}}\right)^{1/\lambda} \quad (1)$$

where x is the price of the auctioned item, $x_i^{\max}$ is the maximum price that this agent is willing to pay for g, and $x_i^{\min}$ is the initial bid that this buyer offers according to the deployed bidding strategy. The parameter λ determines the convexity of the valuation function. Different valuation functions can be obtained for different values of λ. Some of those are shown in Figure 4.

Regarding the remaining-time issue, we assume that all buyers exhibit a rational bidding behaviour according to which they are willing to bid more often as time approaches the end of the auction. In order to model this behaviour, we introduce the probability P_i, with which a buyer submits a new bid at time t. The distribution function of P_i is:

$$P_i(t) = (t/t_{max})^{1/\gamma} \quad (2)$$

where γ is a real number used to regulate the convexity of the curve. Three instances of the bidding probability for γ=2, γ=1, and γ=0.5, are illustrated in the graph of Figure 5. They represent three different time-dependent bidding behaviours that can be characterized as *anxious*, *aggressive* and *conservative*, respectively.

Within the context of the above, the overall valuation function υ_i (x, t) for buyer b_i can be modelled as the product of two components: one, which is resource-dependent, and another, which is time-dependent. Thus, the overall valuation function becomes:

$$\upsilon_i(x, t) = \upsilon_i^p(x) P_i(t) \quad (3)$$

where $\upsilon_i^p(x)$ and $P_i(t)$ are the valuation function of price and the probability of bidding at a given time, given by equations (1) and (2), respectively.

Bidding tactics

A bidding tactic for agent b_i is a function that calculates the price of a potential new offer a buyer intends to submit. In our implementation we use the time-dependent tactic functions $a_i(t)$ introduced by (Faratin *et al.* 2000) for a single attribute. In this case, the price x of a new bid at time $t \leqslant t_{max}$ is given by equation 4:

$$x = x_i^{\min} + a_i(t)\left(x_i^{\max} - x_i^{\min}\right) \quad (4)$$

where $x_i^{\min}$, $x_i^{\max}$ are the same parameters that appear in equation (1). In general, polynomial and exponential functions $a_i(t)$, parameterized by a value δ, which determines the convexity degree, can produce an infinite number of such functions. Any tactic must adhere to the constraints imposed for $a_i(t)$. These are $a_i(0)=k_i$, $a_i(t_{max})=1$ and $0 \leqslant a_i(t) \leqslant 1$. In our implementation we use the polynomial tactic given by:

$$a_i(t) = k_i + (1 - k_i)\left(\frac{\min(t, t_{\max})}{t_{\max}}\right)^{1/\delta} \quad (5)$$

In equation (5), k_i is a constant that determines the price of the bid to be submitted at the first offer by agent b_i. The value of δ determines the bidding behaviour of each buyer. If δ has a value much lower than one, then the corresponding buyer adopts a conservative behaviour (buyer bids more often as the deadline t_{max} is approaching), while for values of δ much greater than one, the buyer adopts a more anxious behaviour (starts bidding from the beginning of the auction).

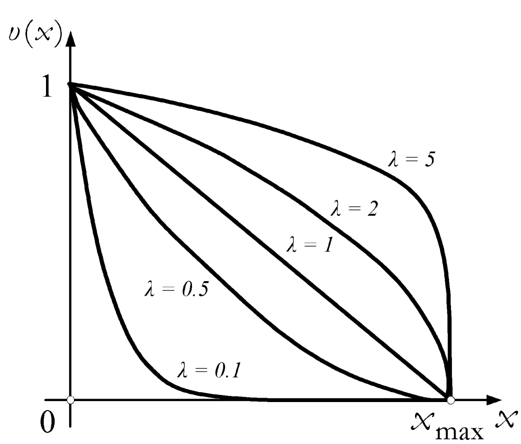

Figure 4. Possible valuation functions for a buyer of the marketplace with respect to the price of the item to be auctioned for several values of the parameter λ

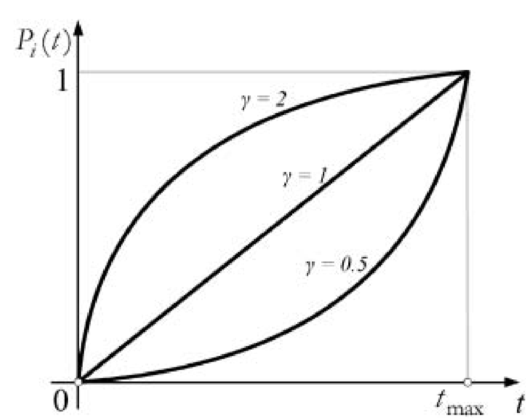

Figure 5. Distributions of the probability of bidding with respect to time t, for three values of the parameter γ: 2 (anxious), 1 (aggressive) and 0.5 (conservative behaviour)

A simple strategy

Some elements of the strategy employed by the buyers have already been mentioned in the description of a *transaction* between a buyer and a seller, stated in definition 2. In particular, the bidding strategy employed by the buyers can be summarized as follows: at an arbitrary time t, each buyer, using equations (4) and (5), calculates a new bid x to potentially submit in the next time interval t+1. Then the following heuristic rule is applied for buyers:

R_1: **IF** $x>\beta$ **AND** $\upsilon_i(x, t+1)>\upsilon_i(\beta, t)$ **THEN** bid x
ELSE no bid (wait)

This simple rule states that if one buyer's bidding tactic proposes a bid greater than β, this bid will be eventually submitted only if it is also profitable for the buyer. Otherwise, the buyer submits no bid. Different values of the parameter δ ensure a great variety of bidding tactics, hence many different strategies (theoretically infinite) can exist among the competitive agents. By heuristic rule R_1 it is also derived that, when the outstanding bid is greater than, or equal to the one currently intended to be submitted by the buyer, it is more profitable for the buyer not to submit a new bid. In this case he waits until the valuation of the outstanding bid substantially increases.

In general, heuristic rule R_1 imposes the necessary constraints to ensure the deployment of a rational bidding behaviour from the buyer's side.

BIDDING WITH FORECASTING

The bidding mechanism introduced so far forms a type of essential bidding agents, which are referred to as 'Simple Agents' (SA). The intermediate information revealed as the auction progresses can be used for short-term forecasting of the next bid. In our approach we use auto-regression modelling, a time-series analysis technique suitable for forecasting. The history of submitted bids forms a time-variant sequence that can be used to estimate the next member using linear projection on a certain number of previous samples. We chose the auto-regression scheme because it is easy to implement and exhibits a low computational cost. In what follows we give an overview of the forecasting technique, which forms the basis of the bidding mechanism. After the predicted bid is calculated, it is inserted into a new reasoning component, which has more rules than the one of SA. This forms an improved knowledge base, incorporated by a new type of bidding agents, characterized as 'Forecasting Agents' (FA). The goal of our experiments is to prove that, although these two types of agents use the same valuation and next-bid calculation functions, described by equations (4) and (5), respectively, FA bidders outperform SA ones.

The prediction mechanism

The prediction algorithm constitutes the core component of the presented reasoning mechanism. Given a sequence of m samples, whose values vary in a timely fashion, the algorithm estimates the value of the m+1 element. Storing the values of the outstanding bid for each t we produce a sequence of time-dependent samples. Based on this, a forecasting schema is deployed, whose goal is to efficiently predict the next outstanding bid. This new parameter is then involved in a new set of heuristic rules (described later on) that introduce an updated version of the previous simple strategy. The prediction algorithm used in our approach is based on a linear estimation of the next member of a weakly stationary sequence of samples using additive noise (Hamilton 1994).

Let $\beta_t, \beta_{t-1}, \ldots, \beta_{t-m+1}$ be a sequence of the previous m values of the outstanding bid. In order to estimate the outstanding bid $\hat{\beta}$ at time $t=m+1$, we can use a linear projection given by equation (6):

$$\hat{\beta}=\sum_{t=0}^{m}\varphi_t\beta_t+\varepsilon_t \tag{6}$$

where ε_t is a white-noise series. Equation (6) describes an *Auto-Regressive* (AR) model of order m. The design of the AR model mainly involves the decision of the value of m and the calculation of the coefficients ϕ_t. In our implementation we selected the AR model of order two (m=2), denoted as AR(2). It is clearly stated (Hamilton 1994) that AR models of higher order are difficult to manipulate as they increase the calculation complexity and the forecast error. In order for equation (6) to represent a valid next bid estimator, the sequence β_t must be properly modified. A method commonly used for this type of data processing for forecasting purposes is the one introduced by Box and Jenkins (Box *et al.* 1994). More details on the attributes of the estimator of equation (6) can be also found in (Hamilton 1994).

Improving the bidding strategy using forecasting

Using the estimator given by equation (6) we introduce a new forecasting strategy, based on a number of heuristic rules. Any buyer that adopts the new strategy first calculates his potential next bid x using equations (4) and (5) and then makes a decision on bidding from the rule-base described in Table 1. Among the

Table 1. The improved rule-base of FA buyers

R_2:	IF $x>\beta$ AND $\hat{\beta}\geqslant x$ THEN $x'=x+\left(\hat{\beta}-x\right)\times a_i(t+1)$ go to R_4
R_3:	ELSE IF $x>\beta$ AND $\hat{\beta}<x$ THEN $x'=\hat{\beta}+\left(x-\hat{\beta}\right)\times a_i(t+1)$ go to R_4
R_4:	IF υ_i (x', t+1)>υ_i (β, t) THEN bid x' ELSE do not bid
R_5:	IF $x<\beta$ THEN do not bid
R_6:	IF $\hat{\beta}<\beta$ THEN go to R_1

possible relations of the three parameters x, $\hat{\beta}$ and β we ignore the case where $\hat{\beta}<\beta$, because it results to an invalid estimation $\hat{\beta}$. The other valid relations holding between the three parameters are $\beta<x\leqslant\hat{\beta}$, $\beta\leqslant\hat{\beta}<x$, and $x\leqslant\beta\leqslant\hat{\beta}$. In the first two cases, the rules presented in Table 1 are applied. In the third case, the buyer chooses not to bid. The heuristic rules R_2–R_6 can handle all possible cases.

Rule R_2 states that if the predicted value $\hat{\beta}$ of the outstanding bid is greater than or equal to the bid x the buyer intends to submit, x is updated to a slightly increased value in proportion to the difference $\hat{\beta}-x$. Recall from equation (5) that the $a_i(t)$ function returns values within the interval [0, 1]. This implies that for any new bid x' proposed by R_2 it is true that $x\leqslant x'\leqslant\hat{\beta}$. At this point, rule R_4 is used to evaluate x' against the current outstanding bid β, according to this buyer's valuation function υ_i. Hence, the buyer finally submits x' only if its valuation is greater than the one of the current outstanding bid. It seems that R_2 includes an element of risk, since it dictates a bid greater than what the buyer intended to do. However, the application of R_4 ensures that the buyer keeps its profit at a desired level, and does not eventually bid unless it is profitable to do so.

In the case where rule R_3 is applied, the value of the next outstanding bid is estimated to be less than x. At this point it would be profitable for the buyer to submit x. However, R_3 suggests the submission of x', which is slightly lower than x. Indeed, R_3 and equation (5) imply that: $\hat{\beta}\leqslant x'\leqslant x$, hence submitting x' is a more profitable decision. Then, heuristic rule R_4 is also applied in this case to evaluate x' and make a decision on its submission. The presence of R_4 in the rule-base is mandatory for a buyer to guarantee rationality.

Heuristic rule R_5 is applied when x is lower than the current outstanding bid. In this situation, buyer decides not to bid, because this decision is more profitable.

Finally, rule R_6 is included in the rule-base to handle circumstances where the estimation is erroneous. These occur when $\hat{\beta}<\beta$. In this case, forecasting is disregarded and the decision on the next bid is handled by rule R_1.

EMPIRICAL EVALUATION

In order to test the efficiency of the proposed mechanism, we have developed a simulated e-marketplace environment and a standard evaluation procedure. In our experiments, we instantiate several agents employing different bidding-tactics and valuation functions. In the beginning of the experimental procedure, we measure the accuracy of the prediction mechanism described by equation (6). We next focus our experiments on monitoring and benchmarking the behaviours of two agent types, formerly defined as SA and FA. The first one possesses a simple strategy described by heuristic rule R_1 and equations (1) through (5). The second agent type employs a more sophisticated strategy, which also includes rules R_2 to R_6.

Evaluation procedure

Our evaluation procedure is mainly interested in two issues:

1. to ensure that the prediction mechanism described by equation (6) can provide a reliable estimation, that is, one with a low error, and
2. to identify any improvements in the bidders' performance introduced in the reasoning component of the FA.

In the first set of experiments, the average of the calculated prediction error is measured with respect to: a) the number of participating agents in the auction, which determines the demand of a particular item g; and b) the auction duration, t_{max}. In this case, we initiate several auctions for one item g, with many agents whose number varies from two to 100. We use as a metric for the prediction quality the normalized squared error: $e=\left|\left(\hat{\beta}-\beta\right)/\beta\right|^2$. The bidding tactics and the valuation functions of an agent that enters an auction are chosen randomly. In particular, we choose for all buyers different values of the parameters λ, γ, and δ, appearing in equations (1), (2) and (5) respectively, thus creating different bidding behaviours.

In order to test the efficiency of the proposed strategy, we introduce a second set of experiments where the performance of each agent type is measured in terms of its average satisfaction: $(1/N_K)\sum_i^{N_K}\upsilon_i$ where N_K is the number of agents of type K and υ_i is the satisfaction degree of agent i. If the latter wins the auction at price x_w, then υ_i is calculated from equation (1) for λ=1, x_{min}=0, x=x_w, otherwise it is taken equal to zero. In these experiments a set of FA bidders is compared to another set of SA bidders with similar configuration parameters determined by the value of γ. We considered three cases: γ=2, γ=1, and γ=0.5, corresponding to the three time-dependent bidding behaviours, previously characterized as anxious, aggressive, and conservative. The parameters in equations (1) and (5) are assigned the values: $\lambda=10$, $\delta=1$, $k_i=0.01$, while the duration of an auction increases linearly from 10 to 100 time units. Each of the aforementioned configurations is applied on 90 experiments, resulting to a total of 270 different auctions. The results are presented in Figures 6 and 7.

Results

In this section, selected experimental results are presented and discussed in order to demonstrate improvement in the decision mechanism of buyers who adopt our methodology. The first set of experiments aims at testing the accuracy of the prediction mechanism. This is measured by having an auction executed for 100 times. In each auction a new agent enters the marketplace, causing an increase on the demand. Figure 6(a) shows the average prediction error with respect to the demand. In this figure, it is shown experimentally that the demand has no significant influence on the prediction error. Indeed, from figure 6(a), it is shown that, as the number of agents increases, the average prediction error behaves as a linear constant function, for sizes of agent population greater than 10. For a small number of agents, the average prediction error has a slight variation. Another very interesting point is that the error is remarkably low. The second experiment for testing the forecasting accuracy is performed with respect to the

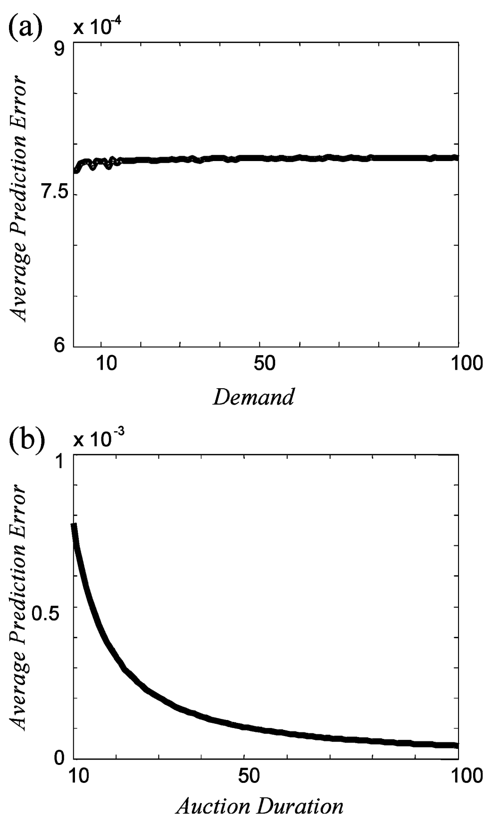

Figure 6. The average prediction error with respect to (a) demand and (b) duration of an auction

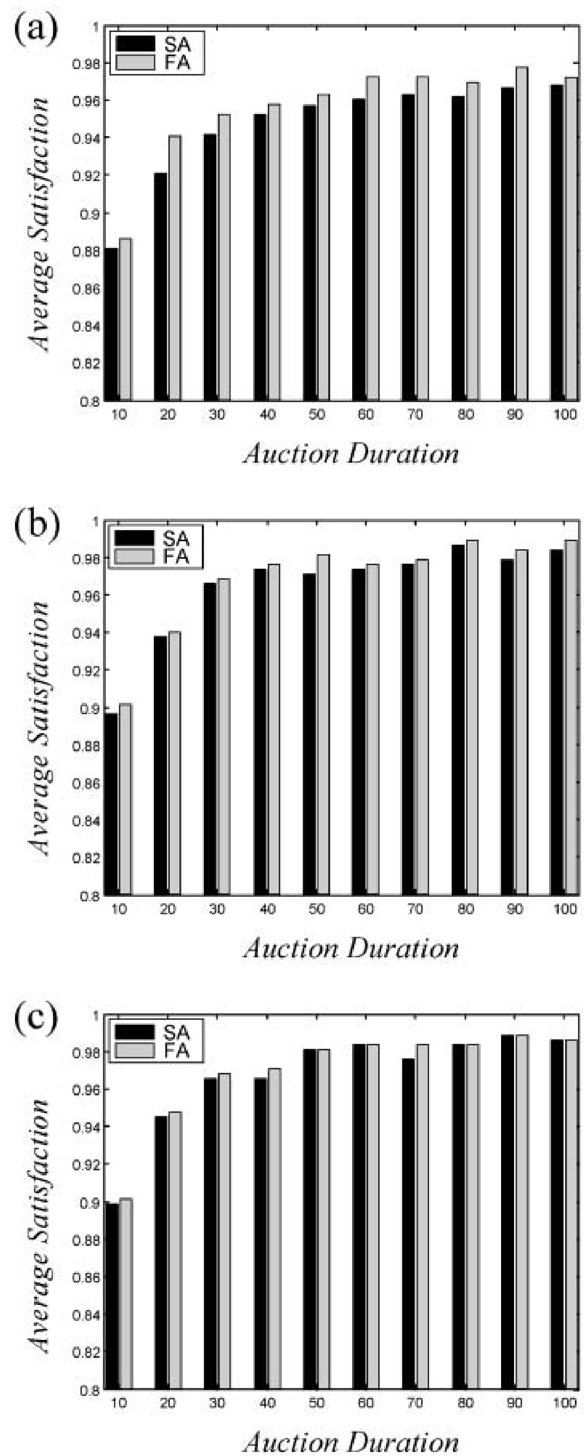

Figure 7. Performance of FA and SA agent types for three different bidding behaviours with respect to the duration of the auction: (a) anxious, (b) aggressive and (c) conservative

duration $t_{\max}$ of the auction. The calculated average prediction error is illustrated in figure 6(b). Apparently, the prediction error decreases exponentially, as the auction duration increases. This result demonstrates that our forecasting estimator is being improved over the auction time, indicating that it behaves more efficiently in long auctions.

In the second set of experiments, conducted in order to measure the FA performance, we run a particular number of auctions for the three classes of bidding behaviours, keeping their attributes constant, while increasing the duration of each auction. First, we examine the behaviour of an anxious FA, which participates in the same environment with a group of anxious SAs. This implies that in the experimental environment the demand does not change. The average satisfaction of each agent type with respect to the auction duration is shown in Figure 7(a). In Figures 7(b) and 7(c) we repeat the same experiment with a FA and a SA that deploy aggressive and conservative behaviours, respectively. In all three cases the FA consistently outperforms the SA. The FA bidders also increase their average utility as the duration of the auction becomes longer. The results of Figures 7(a)–(c) indicate that the adoption of forecasting results in improved bidding behaviours.

CONCLUSIONS AND FUTURE WORK

This paper has presented a bidding strategy, which is based on the prediction of the next outstanding bid in an e-marketplace with competitive autonomous agents. For experimenting on the proposed method we have implemented an electronic auction environment with one auctioneer and many buyer agents, acting in a competitive way, in order to achieve the best buying price for one auctioned item. We have described formally the discrete bidding mechanisms that the agents of the auction environment deploy and we have also introduced FA, an agent that deploys the proposed, enhanced, forecasting mechanism. The latter is based on a number of heuristic rules, which are properly outlined and explained. In order to evaluate the performance of this bidding mechanism, two types of experiments have been conducted: a) for certain agent parameters, we have monitored the prediction error of the forecasting model with respect to auction duration time and with increasing agent demand, and b) for the same agent parameters, we have monitored the profit of the forecasting model, with respect to auction duration. Results in both cases appear interesting and promising. The prediction error proves to be dependent on the duration of the auction but not on the demand. Measuring profit against time shows that the FA performance increases sufficiently.

Further research work will be focused on two directions: a) the incorporation of fuzzy criteria in the rules of the FA knowledge-base, and b) the incorporation of an inductive reasoning mechanism, by performing data mining techniques on historical data. As far as the former enhancement is concerned, related research literature has indicated that fuzzy rules can provide better understanding and manipulation of bidding events in an auction. Concerning the second potential improvement, the processing of sequential historical auction data by the use of data mining techniques (that is trend analysis, sequential pattern mining) may reveal hidden similarities within the bidding behaviours, that can be used by the FA in a more successful forecasting, and consequently, bidding behaviour.

References

Beam, C. and Segev, A. (1997) 'Automated Negotiations: A Survey of the State of the Art', *Wirtschaftsinformatik* 39(3): 263–8.

Bichler, M. (2000) 'A Roadmap to Auction-based Negotiation Protocols for Electronic Commerce', *Proceedings of the 33th Hawaii International Conference on System Sciences*, Washington, DC: IEEE Computer Society.

Box, G., Jenkins, G. M. and Reinsel, G. (1994) *Time Series Analysis, Forecasting and Control*, 3rd edn, Englewood Clifs, NJ: Prentice Hall.

Faratin, P., Sierra, C. and Jennings, N. (2000) 'Negotiation Decision Functions for Autonomous Agents', *Inernatonal Journal of Robotics and Autonomous Systems* 24(3–4): 159–82.

Greenwald, A. and Stone, P. (2001) 'Autonomous Bidding Agents in the Trading Agent Competition', *IEEE Internet Computing* 5(2): 52–60.

Hamilton, J. D. (1994) *Time Series Analysis*, Princeton University Press, 56–8.

He, M., Leung, H.-F. and Jennings, N. (2002) 'A Fuzzy Logic Based Bidding Strategy for Autonomous Agents in Continuous Double Auctions', *IEEE Transactions on Knowledge and Data Engineering* 15(6): 1345–63.

He, M., Jennings, N. and Leung, H.-F. (2003) 'On Agent-mediated Electronic Commerce', *IEEE Transactions on Knowledge and Data Engineering* 15(4): 985–1003.

Kurbel, K. and Loutchko, I. (2002) 'Multi-agent Negotiation Under Time Constraints on an Agent-based Marketplace for Personnel Acquisition', in T. Dittmar, B. Franzzyk and R. Hofmann (eds) *Proceedings of the 3rd International Symposium on Multi-Agent Systems, Large Complex Systems, and E-Business* (MALCEB 2002), Erfurt, 566–79.

McAffee, P. R. and McMillan, J. (1987) 'Auctions and Bidding', *Journal of Economic Literature* 25: 639–738.

Milgrom, P. R. (1989) 'Auctions and Bidding: A Primer', *Journal of Economic Perspectives* 3: 3–22.

Sandholm, T. (1999) 'Distributed Rational Decision Making', in G. Weiss (ed.) *Multiagent Systems*, The MIT Press, 201–58.

Vickrey, W. (1961) 'Counterspeculation, Auctions, and Competitive Sealed Tenders', *Journal of Finance* 16: 8–37.

Wellman, M., Greenwald, A., Stone, P. and Wurman, P. R. (2003a) 'The 2001 Trading Agent Competition', *Electronic Markets* 13: 4–12.

Wellman, M., Reeves, D., Lochner, K. and Vorobeychik, Y. (2003b) 'Price Prediction in a Trading Agent Competition', *Journal of Artificial Intelligence Research* 19: 901–18.