Possibilistic-Based Design of Bidding Strategies in Electronic Auctions

Pere Garcia[1], **Eduard Giménez**[1], **Lluís Godo**[1] and **Juan A. Rodríguez-Aguilar**[1]

Abstract. Auction-based electronic commerce is an increasingly interesting domain for developing trading agents competing in multi-agent electronic markets. In this paper we present an attempt towards the construction of such agents, introducing both a formal and a more pragmatical approach to the design of bidding strategies for buyer agents in auction-based tournaments. Our formal view relies on possibilistic-based decision theory as the means of handling possibilistic uncertainty on the consequences of actions (biddings) due to the lack of knowledge about the other agents' behaviour. For practical reasons we propose a two-fold method for decision making that does not require the evaluation of the whole set of alternative actions. This approach utilizes global (market-centered) information in a first step to come up with an initial set of potential bids. This set is subsequently refined in a second step by means of the possibilisitic decision model using individual (rival agent centered) information induced from a memory of cases composing the history of tournaments. The proposed approach has been implemented and it is currently being tested in tournament scenarios defined with FM97.6, a multi-agent testbed for electronic auctions.

1 Introduction

Auction-based electronic commerce appears to be an area where the Web is proving to be better than traditional alternatives due mainly to its highly interactive nature and to the fact that online auctions significantly reduce costs. Thus, online auctions such as Auctionline(www.auctionline.com), Onsale (www.onsale.com), InterAUCTION (www.interauction.com), eBay (www.eBay.com) and many others have proliferated over the Internet.

Thus, auctions must be regarded as an attractive domain for developing shopping agents. Nevertheless, the functionality shown by shopping agents like Roboshopper(www.roboshopper.com), Jango (www.jango.com), or Shopping Explorer(www.shoppingexplorer.com) amounts to looking for products through the Web, conveying the retrieved information to the users, and possibly assisting them in their purchases. A natural extension would be to enpower them with actual trading capabilities. The AuctionBot [9] project is a first attempt in this direction.

Trading within an auction house demands that buyers to decide on an appropriate price on which to bid. But this decision —if rational— should profit from whatever information may be available in the market: participating traders, available goods and their expected re-sale value, historical experience of prices and participants' behaviour, etc. However, richness of information is not the only source of complexity in this domain. The actual conditions for deliberation may be constantly changing and highly uncertain and furthermore, deliberations are significantly time-bounded. Consequently, in such a context, the trading agent's decision-making process may be quite elaborate and thus, designing, building, and tuning trading agents becomes a difficult task.

In [5] we introduced FM97.6, a multi-agent testbed for electronic auctions that intends to provide support to agent developers. FM97.6 permits the definition, activation, and evaluation of experimental game-like trading scenarios called $tournaments$ which define standardized conditions under which agents compete for maximizing their benefits. The problem of choosing a succesful bidding strategy by a trading agent in these type of n-agent auction tournaments is clearly not deterministic and depends on many factors, in particular on the strategies of the other competing agents. Since the knowledge the agent has about the other agents' strategies is usually incomplete, the approach presented in this paper consists in regarding this problem as a decision making problem under uncertainty.

As in any decision problem, the trading agent has to choose a decision, i.e. a bid, among a set of available alternatives, taking into account her preferences on the set of possible consequences in terms of maximising her benefit. In decision problems, given a (finite) set of situations Sit and a (finite) set of consequences X, a non-uncertain decision d is represented by a function $d : Sit \rightarrow X$. Then, in each situation, decisions can be easily ranked using, for instance, a real-valued utility function $u : X \rightarrow \Re$ modelling the preferences between consequences. However, uncertainty may be involved in many different aspects of the decision process, in particular, in a given situation, we may be uncertain on what the consequences of a decision are. Classical approaches to decision making under uncertainty assume that uncertainty is represented by probability distributions. In such a case, if we are in a situation s_0 and we know the probability $P_{s_0,d}(x)$ of each possible consequence $x \in X$ of a decision d, then the global utility of the decision d is usually evaluated as the expected value of u with respect to $P_{s_0,d}$:

$$\mathcal{U}_{s_0}(d) = \Sigma_{x \in X} P_{s_0,d}(x) u(x).$$

This utility is used then to rank decisions, the greater, the better. This kind of approach corresponds to the well-known Expected Utility Theory (EUT) [4], but it presents some problems and paradoxes, basically related to inferring the probabilities. Indeed, in our problem of trading agents the working assumption is that the knowledge the agent has about the other agents' strategies is reduced to a memory (or history) of previous successful biddings. In such a framework, we propose a kind of *case-based reasoning* to observe behaviour in previous similar situations. The uncertainty induced in these kind of processes is due to what extent we may consider two situations similar enough to presume a similar behaviour, and this kind of uncertainty

[1] Artificial Intelligence Research Institute (IIIA), Spanish Council for Scientific Research (CSIC), 08193 Bellaterra, Barcelona, Spain.

ECAI 98. *13th European Conference on Artificial Intelligence*
Edited by Henri Prade
Published in 1998 by John Wiley & Sons, Ltd.

is possibilistic rather than probabilistic. Recently, Dubois and Prade have proposed in [3] a qualitative decision theory which relies on the basic assumption that uncertainty about a decision problem is of *qualitative* nature and representable by *possibility distributions* over the set of consequences X. Given a qualitative[2] uncertainty scale V, a possibility distribution $\pi : X \to V$ provides a plausibility ordering on X in such a way that $\pi(x') \leq \pi(x)$ means that the consequence x is at least as plausible as x'. Based on such uncertainty representation, (qualitative) utility functions are defined and characterised by a set of axioms which can be regarded as a qualitative counterpart of Von Neumann and Morgenstern's axiomatics for the EUT.

In this paper we build on this decision model to design bidding strategies for trading agents. After this introduction, Section 2 briefly describes the auction tournament environment FM97.6. In Section 3, our theoretical decision model is described while Section 4 is devoted to explain how this decision model can be used in a two step, heuristic-based approach for agent's decision making. Finally, we end up with some conclusions and an outline of our future work.

2 Defining Tournament Scenarios

FM96.5 (see [6, 8]) is an electronic auction house based on the traditional *fish market* metaphor as an alternative to other proposals for electronic marketplace architectures [7, 1]. In a highly mimetic way, the workings of FM96.5 involves the concurrency of several scenes governed by the market intermediaries identified in the fish market. Namely, seller agents register their goods with a seller admitter agent, and wait for receiving their earnings from a seller manager until the auctioneer has sold them in the auction room. Buyers, on the other hand, register with a buyer admitter, and bid in the auction room for goods which they pay through a credit line that is set up and updated with a seller manager. The principal scene is the auction itself, in which buyers bid for boxes of fish that are presented by an auctioneer who calls prices in descending order - the *downward bidding protocol*. Buyer and seller agents can trade goods as long as they comply with the *FishMarket institutional* conventions, incarnated in what we call a *remote control*, which constitutes the sole and exclusive means through which a trader agent —be it a software agent or a human trader— interacts with the market institution. A remote control gives a permanent identity to the trader and enforces an *interaction protocol* that establishes what illocutions can be uttered by whom and when.

In order to obtain an auction tournament environment, FM96.5 has been extended with some innovations that turn it into a multi-agent testbed, FM97.6, already described in [5]. FM97.6 permits the definition, activation and evaluation of a wide variety of experimental trading scenarios (from simple toy scenarios to complex real-world scenarios) that we shall refer to as *tournaments*.

A tournament scenario will involve a collection of explicit parameters that characterize an artificial market. Such parameters define the bidding conditions (timing restrictions, increment/decrement steps, publicly available information, etc.), the way goods are identified and brought into the market, the resources buyers may have available, and the conventions under which buyers and sellers are going to be evaluated.

In rest of this section we sketch out our (formal) view of such tournament scenarios by introducing some of the elements composing them[3]. We start by explicitly describing the dynamics of the *downward bidding protocol* governing the main activity within the *FishMarket*:

[Step 1] The auctioneer chooses a good out of a lot of goods that is sorted according to the order in which sellers deliver their goods to the sellers' admitter.

[Step 2] With a chosen good, the auctioneer opens a *bidding round* by quoting offers downward from the good's starting price, previously fixed by the sellers' admitter, as long as these price quotations are above a *reserve price* previously defined by the seller.

[Step 3] Several situations might arise during this round:

- **Bids:** One (or several) buyer(s) submit his/their bids at the current price. If there is only one bid, the good is sold to the bidder. Otherwise, a collision comes about, the good is not sold, and the auctioneer restarts the round at a higher price.
- **No bids:** No buyer submits a bid at the current price. If the reserve price has not been reached yet, the auctioneer quotes a new lower price, otherwise the auctioneer declares the good *withdrawn*[4] and closes the round.

[Step 4] The first three steps repeat until there are no more goods left.

Next, we identify the elements composing a *FishMarket* -like tournament scenario by introducing the notion of *Tournament descriptor*, which intends to encompass all the information characterizing tournament scenarios using the above downward bidding protocol. We define a **Tournament Descriptor** $\mathcal{T}$ as the 6-tuple

$$\mathcal{T} = \langle \Delta_{price}, \mathcal{B}, \mathcal{S}, Cr, \mu, E \rangle$$

where Δ_{price} is the decrement of price between two consecutive quotations uttered by the auctioneer; $\mathcal{B} = \{b_1, \ldots, b_n\}$ is a finite set of identifiers corresponding to all participating buyers, analogously $\mathcal{S}$ for the participating sellers; Cr is the initial endowment of each buyer at the beginning of each auction; $\mu \in \mathcal{M}$ is the tournament mode where $\mathcal{M} = \{$random, automatic, one auction, fish market, $\ldots\}$ is the set of tournament modes; and E is the buyers evaluation function calculated as the cumulative benefits.

It is of worth noticing that a number of experimental tournament scenarios of varying degrees of realism and complexity can be generated by instantiating the definition above[5].

3 Possibilistic-based Decision for Designing Bidding Strategies

This section aims at outlining our proposal for the development of competitive bidding strategies that help a buyer agent succeed in the *FishMarket* tournaments proposed in [5, 8]. Within such settings, an agent's bidding strategy must decide on an appropriate price on which to bid for the good to be auctioned during a given round. We start in the next subsection by summarizing Dubois and Prade's possibilistic-based decision model. The remaining of this section describes how this model is adapted for our trading agents.

3.1 Background on Possibilistic-based Decision

First we introduce some notation and definitions:

[2] In the sense that the ordering is the only that matters.

[3] The interested reader may refer to [5] for a more thorough discussion.

[4] The good is returned to its owner.

[5] [5, 8] provide examples of tournament instantiations.

- Sit will denote the set of situations and $X = \{x_1, \ldots, x_p\}$ will denote a finite set of consequences.
- $(V, \leq)$ will denote a (finite) linear scale of uncertainty where $\inf(V) = 0, \sup(V) = 1$.
- $\Pi(X)$ will denote the set of consistent possibility distributions on X over V, i.e. $\Pi(X) = \{\pi : X \to V \mid \exists x \in X$ such that $\pi(x) = 1\}$.
- $(U, \leq)$ will denote a (finite) linearly ordered scale of preference, or utility, with $\sup(U) = 1$ and $\inf(U) = 0$. Moreover, $n_U : U \to U$ will denote the order reversing involution in U.
- $u : X \to U$ will denote a utility function that assigns to each consequence x of X a preference level $u(x)$ of U.

We understand a decision d as a mapping $d : Sit \to \Pi(X)$ from situations into possibility distributions on consequences, assuming that, although we know which is the current situation (say s_0), we may partially know which are the consequences of a given decision. Thus, in a given situation s_0, every decision $d \in D$ induces a possibility distribution $\pi_d = d(s_0)$: $\pi_d(x)$ measures the plausibility of consequence x after having taken decision d. Therefore, to rank decisions amounts to rank possibility distributions of $\Pi(X)$, and thus we shall be interested in (gobal) utility functions definable over possibility distributions of $\Pi(X)$.

In this framework, Dubois and Prade [5] propose to use two kinds of qualitative utility functions to rank order possibility distributions. The basic underlying idea is the following one. A utility function $u : X \to U$ on the consequences can be regarded as specifying a fuzzy set of *preferred consequences*: the greater is $u(x)$, the more preferred is the consequence x and the more x belongs to the (fuzzy) set of preferred consequences. On the other hand, a possibility distribution $\pi : X \to V$ specifies which consequences are plausible: the greater $\pi(x)$, the more plausible is the consequence x. Therefore, a pessimistic (or conservative) criterion is to look for those π's which, at some extent, make hardly plausible all the bad consequences. On the contrary, an optimistic criterion is to look for those π's that, also to some extent, make plausible some of the good consequences.

However, to define such criteria an assumption of *commensurateness* between the plausibility scale V and the preference scale U has to be made. In [3] it is assumed to exist an order preserving mapping $h : V \to U$, such that $h(1) = 1$ and $h(0) = 0$, relating both scales. Then, given a utility function $u : X \to U$ the following pessimistic and optimistic qualitative utilities for each decision d are defined respectively:

$$QU^-(d \mid u) = \min_{x \in X} \max(n_U(\pi_d^*(x)), u(x))$$
$$QU^+(d \mid u) = \max_{x \in X} \min(\pi_d^*(x), u(x))$$

where $\pi_d^*(x) = h(\pi_d(x))$. We notice that, $QU^-(d \mid u)$ and $QU^+(d \mid u)$ are the Sugeno integrals of the utility function u with respect to the necessity and possibility measures induced by the distribution π_d^* [3]. Moreover, when π_d^* is crisp (i.e. either $\pi_d^*(x) = 1$ or $\pi_d^*(x) = 0$ for any $x \in X$), QU^- and QU^+ turn out to correspond the well-known *maximin* and *maximax* decision criteria.

Although Dubois and Prade's decision model explicitely assumes to deal with finite scales of uncertainty (V) and utility (U), for conceptual simplicity, and also for implementation reasons (e.g. testing flexibility), we shall take from now on the unit interval [0,1] to be both the uncertainty and utility scale, i.e. $V = U = [0, 1]$. It will be then possible to consider other expressions (of a more quantitative nature) for the utility besides the pure extreme qualitative utilities QU^- and QU^+. For instance one could also think of the following numerical utility:

$$\mathcal{U}(d \mid u) = \frac{\Sigma_{x \in X} \pi_d(x).u(x)}{\Sigma_{x \in X} \pi_d(x)},$$

accounting also for a degree of inclusion of the set plausible consequences into the set of good consequences. Notice that, if both π_d and u define crisp subsets of consequences, then $\mathcal{U}(d \mid u)$ is nothing but the relative cardinality of π_d inside u (w.r.t π_d), and obviously $\mathcal{U}(d \mid u) = 1$ only if $\pi_d \subseteq u$.

3.2 Instantiating the decision model for our trading agents

In order to apply the possibilistic decision model first we have to identify the variables involved in the decision problem of our interest.

We model market situations faced by our agent, denoted hereafter b_0, as vectors of features

$$s = (r, a, g, p_s, p_r, \overline{\kappa}, \overline{E}, R)$$

characterizing round r of auction a such that g is the good to be auctioned, p_s is its starting price, p_r is its resale price, $\overline{\kappa}$ is the vector of current credits (κ_i is the current credit of buyer b_i), $\overline{E}$ is the vector of scores (E_i is the score of buyer b_i), and R is the number of rounds left.

The decision set $\mathcal{D}$ will consist of the set of allowed bids our agent b_0 can perform. In situation s_0, we shall have $\mathcal{D} = \{Bid(p) \mid p = p_s^0 - m.\Delta_{price}, m \in \mathbb{N}, p_r^0 \leq p \leq Cr(b_0)\}$, where p_s^0 and p_r^0 are the starting and reserve prices in situation s_0, and $Bid(p)$ means that the agent makes a bid at price p.

In each round, either our agent (b_0) wins, or buyer b_1 wins, ... , or buyer b_n wins, paying a given price p for the auctioned good. Therefore, the set X of outcomes (or consequences) in each round is the set $X = \{(Win(b_i), p) \mid i = 0, \ldots, n\}$, where $x = (Win(b_i), p)$ means that the winner of the round is buyer b_i paying a price p.

So, according to the decision model introduced in the previous subsection, given a current market situation s_0, it remains to assess, for each possible decision (bid) $d \in \mathcal{D}$, which are the possibility and utility values $\pi_d(x)$ and $u(x)$, for all $x \in X$, to be able to calculate a global utility for each d (using either QU^-, QU^+, or $\mathcal{U}$). This is done in the next subsections.

3.3 Generating possibility distributions from cases

The key idea is observing the behaviour of each buyer agent in previous similar market situations. Then the uncertainty on the behaviour of an agent in front of a given market situation may be estimated as possibility degrees, in terms of the extent we may consider the current situation similar to those where that agent exhibited that behaviour.

More formally, suppose we have stored a set of precedent decision problems in past rounds and auctions as triples $(s, Win(b), p)$, which we will call *cases*, where s is the situation descriptor, $Win(b)$ denotes that buyer b won the round, and p was the price paid by b for the auctioned good. M will denote the global memory of cases the agent may access.

Given a present situation $s_0 = (r^0, a^0, g^0, p_s^0, p_r^0, \overline{\kappa}^0, \overline{E}^0, R^0)$, for each possible bidding $d = Bid(p_d)$ our agent b_0 may do, she has to evaluate at what extent each buyer (including herself) may win the round at a given price, obviously higher than p_d. This is done by looking for cases at the memory M with similar situation descriptors and similar prices. Namely, for each buyer b_i and each price p, we consider those cases $(s, Win(b_i), p')$ in M and apply the following heuristic principle[6]: "The *more similar* s_0 is to s and p is to p', the *more possible* b_i, paying a price p, will be the winner in s_0". This can be interpreted as the following constraint on possibility values:

$$\pi_{s_0}(Win(b_i), p) \geq \min(\mu_{\tilde{s}}(s_0), \mu_{\tilde{p'}}(p)),$$

where $\mu_{\tilde{s}}$ and $\mu_{\tilde{p'}}$ denote the membership functions of the fuzzy sets $\tilde{s}$ and $\tilde{p'}$ of situations and prices similar to s and p' respectively. These constraints turn into

$$\pi_{s_0}(Win(b_i), p) = \max_{(s, Win(b_i), p') \in M} \min(\mu_{\tilde{s}}(s_0), \mu_{\tilde{p'}}(p))$$

when all the cases in the memory are considered. But of course, there is also a hard constraint that has to be taken into account: b_i $(i \neq 0)$ cannot win paying a price $p < p_d$ if our agent b_0 makes the bid at price p_d. Altogether leads to the following estimate of the possibility degrees of each outcome $(Win(b_i), p)$, for $i \neq 0$:

$$\pi_{s_0, p_d}(Win(b_i, p)) = \begin{cases} \pi_{s_0}(Win(b_i), p), & p_\alpha \geq p \geq p_d \\ 0, & otherwise \end{cases}$$

Finally, to estimate of the plausibility that our agent b_0 wins the round with the bid p_d we look at M for those cases in which in similar situations the price paid was not higher than p_d:

$$\pi_{s_0, p_d}(Win(b_0, p_d)) = \max_{(b,p) \in X} \min(\pi_{s_0}(b, p), \mu_{\tilde{<p_d}}(p)),$$

and $\pi_{s_0, p_d}(Win(b_0, p)) = 0$ for $p \neq p_d$, where $\tilde{<p_d}$ denotes a fuzzy set of prices not higher than p_d.

Last, but not least, there is the question of how to define the fuzzy sets $\tilde{s}$ and $\tilde{p}$, for some situation s and some price p. The idea is to define $\mu_{\tilde{s}}(s_0) = S(s, s_0)$ where $S : Sit \times Sit \to [0, 1]$ is a fuzzy relation on the set of situations accounting for a notion of proximity or similarity[7], and analogously for the prices.

3.4 Assessing utilities

The last elements which remain to specify in our decision problem are, for a given a situation s_0 and a current bidding $d = Bid(p)$, the utility values $u(Win(b_i))$ the agent b_0 assesses at the fact that buyer b_i can be the potential winner of the round. For this purpose, we shall consider consider two alternative utility functions.

The first one aims at modelling an agent that prefers to wait and see when he is ahead, whereas he becomes more and more agressive when he lags behind in order to reach the first position in the tournament. It is based on following scoring function:

$$f(b_i, s_0, p) = \begin{cases} k \cdot t, & \text{if } k \leq 0 \\ k \cdot t^{-1}, & \text{otherwise} \end{cases}$$

where $k = ((\max_{j \neq i} E_j^0) - E_i^0)$ and $t = (R^0 - 1)/(\max(\kappa_i^0 - p, 1) \cdot (p_r^0 - p))$, being p_r^0 the resale price. We assume that $p_r^0 - p \geq 0$, and $\kappa_i^0 - p \geq 0$ i.e., buyers only consider bids that can improve their score, and they have enough credit to submit the bid d. In f the factors $(Max_{j \neq i} E_j^0) - E_i^0$, $p_r^0 - p$, and $\frac{R^0 - 1}{Max(\kappa_i^0 - p, 1)}$ stand for the position of buyer b_i with respect to the other buyers in the ranking of scores, the benefit and estimated cost of winning the round, respectively. Finally, the utility function that we propose in this approach is:

$$u(Win(b_i), p) = \begin{cases} r(f(b_0, s_0, p)) & \text{if } i = 0 \\ r(-f(b_i, s_0, p)) & \text{otherwise} \end{cases}$$

where r is a linear scaling function which makes u fall into [0,1].

Another point of view for defining a utility function could be the following one based on considering the future benefit of an agent as a random variable. Namely, if in a given situation s_0 the outcome $(Win(b_i), p)$ happens, then the next situation is determined since we can update the vector $\overline{E}$ of benefits and the vector $\overline{\kappa}$ of the remaining credits. Thus the utility of the consequence $(Win(b_i), p)$ should express how good is for us to start up the following round with such news descriptor values in terms of the estimate of the final score E_0^F for our agent b_0 w.r.t. the score of the other agents. For this purpose let us assume we are in round j of an auction with a total of N rounds. Then we, for each agent b_i, the final benefit can be computed as

$$E_i^F = E_i^j + \Sigma_{n=j+1}^N benefit_i^n$$

, where $benefit_i^n$ denotes the benefit of agent b_i in round n. We shall consider that $benefit_i^n$ as a value of a random variable describing the benefit that b_i can get at the round$_n$. On this random variable we make two working hypothesis: the first one is that it is proportional to the resale price, that is, $benefit_i^n = p_r \cdot x_i^n$; the second is that it is proportional to the current credit, that is, $benefit_i^n = \kappa_i^n \cdot y_i^n$. This provides two methods to estimate both the mean an standard deviation of the random variables x_i^n and y_i^n from past cases. Finally, since the above variable E_i^F is a sum of a constant with the variables $benefit_i^n$ (indeoendence is assumed) we can approximate its probability distribution $\mathcal{P}$ by the corresponding normal distribution. Since winning an auction means that $E_0^F > E_i^F, \forall i \neq 0$ we can define the utility of any consequence as

$$u(Win(b_i, p)) = \mathcal{P}(E_0^F > \max_{i \neq 0} E_i^F).$$

4 An Heuristic Approach for the Development of Competitive Bidding Strategies

The outright use of the possibilistic-based decision mechanism described above appears to be prohibitively expensive. Therefore, when facing the design of pragmatic bidding strategies we must attempt to propose flexible heuristic guidelines that prevent the agent from evaluating the whole set of alternative actions so that its deliberation process constrains to time and resource-boundedness. In what follows we propose a two-folded approach for decision making. In the first step a set of potential bids (as a subset of the whole set of possible bids) is selected according to the general trend of the markert. Then, the second step consists of selecting the best bid for the agent according to the possibilistic decision model previously described.

We describe next how the first step is put into practice. Let s_0 be the current situation and M the available memory of cases. s_0, a first Assuming the principle that "*similar market situations* usually lead to *similar sale prices* of the good", the idea is to take advantge of the interpolation mechanism implicit in the fuzzy case-based reasoning model proposed in [**?**]. This amounts to consider, for each case[8] $(s, p) \in M$, a gradual fuzzy rule "*If* Σ is $\tilde{s}$ *then* Υ is $\tilde{p}$", where again $\tilde{s}$ and $\tilde{p}$ stand for the fuzzy set of situations similar to s and

[6] A similar principle has been recently considered in [2] in a framework of fuzzy case-based reasoning.

[7] See for instance [**?**] for details about fuzzy similarity relations.

[8] Since we are only interested in the situation descriptor and the sale price we ommit in this section the buyer's identifier in the cases.

the fuzzy set of prices similar to p respectively; Σ and Υ are variables ranging over situations and prices resp. (**Caution**: the fuzzy set $\tilde{s}$ may be different from the fuzzy set with the same name appearing in subsection 3.3, since the criteria used to define how similar are situations may change from one purpose to another.) Such a fuzzy rule, together with the input situation s_0 gives the following fuzzy set $pbid$ of prices (see [**?**]) for details about the semantics of fuzzy gradual rules):

$$\mu_{pbid}(p') = I(\mu_{\tilde{s}}(s_0), \mu_{\tilde{p}}(p')).$$

Here I is residuated many-valued implication[9], and assuming to have defined similarity functions $S^* : Sit \times Sit \rightarrow [0,1]$ and $T^* : prices \times prices \rightarrow [0,1]$, we may define $\mu_{\tilde{s}}(s_0) = S^*(s, s_0)$ and $\mu_{\tilde{p}}(p') = T^*(p, p')$. Considering all the cases in the memory M we come up with the following (fuzzy) set of potential bids:

$$\mu_{pbid}(p') = \min_{(s,p)\in M} I(S^*(s, s_0), T^*(p, p')).$$

Finally the set $\hat{B}_\alpha$ of candidate bids can be selected to be those having a membership degree to $pbid$ above a certain value $\alpha > 0$, i.e. $\hat{B}_\alpha = \{p' \mid \mu_{pbid}(p') \geq \alpha\}$. It can be checked that $\hat{B}_\alpha = \bigcap_{(s,p)\in M}\{p' \mid T^*(p,p') \geq S^*(s, s_0) \otimes \alpha\}$, where $\otimes$ is the t-norm whose residuum is I.

Therefore, instead of considering the whole set of alternative bids, the proposed decision-making process will only evaluate those bids within $\hat{B}_\alpha$. The next algorithm summarizes the process of selecting a bid out of $\hat{B}_\alpha$ using a utility function U.

```
Function Bid_Selection (M, B̂_α, s_0, U))
  ∀ candidate bid p_d ∈ B̂_α
      ∀ buyer b ∈ 𝓑 such that b ≠ b_0
          Retrieve all cases c = (s, b, p_s) ∈ M
          Identify the most similar case c* = (s*, b, p*_s)
          Consequence x_b := (Win(b), p)
          if p ≥ p_d
              then π(x_b) = min(S_1(s_0, s*), S_2(p, p*_s))
              else π(x_b) = 0
        u(x_b, p) = r(−f(b, s_0, p))
      end for
      Retrieve all cases c = (s, b, p_s) ∈ M such that p_s ≤ p
      Identify the most similar case c* = (s*, b*, p*_s)
      π(x_b0, p_d) = S(s_0, s*), π(x_b0, p) = 0, for p ≠ p_d
      u(x_b0, p) = r(f(b_0, s_0, p))
      Calculate U(bid(b_0, p))
  end for
  return p = arg max{U(bid(b_0, p')) | p' ∈ B̂}
end function
```

5 Conclusions and Future Work

We have described a possibilistic-based decision method that attemps at modelling buyer agents' behaviour in electronic auction tournaments.

We have implemented a buyer agent in Java that interacts with the market with the aid of the remote control agent intermediary provided by FM97.6. This agent spawns its remote control as a children process and communicates with it following the communication protocol described in [5] As to the case memory, the cases belonging to the current tournament are kept in the agent's memory whereas cases belonging to past tournaments are kept in a *mySQL* database accessed by the agent through JDBC. In addition to this, our agent has been endowed with libraries of similarity and utility functions in order to empirically determine which similarity and utility functions turn out to be more appropriate for each tournament scenario.

As to our future work, this shall focus on several directions. The most obvious one is to complete the empirical evaluation of our proposal that surely will provide us with valuable feedback to improve our current model. Secondly, the conception of bidding strategies capable of facing other auction protocols since goods can be traded under any auction protocol in FM97.6.

[9] A residuated many-valued implication is a binary opereation in [0, 1] of the form $I(x, y) = \sup\{z \in [0,1] \mid x \otimes z \leq y\}$, where $\otimes$ is a t-norm, i.e., a binary, non-decreasing, associative and commutative operation in [0,1] such that $x \otimes 1 = x$ and $x \otimes 0 = 0$.

ACKNOWLEDGEMENTS

This work has been partially supported by the Spanish CICYT project SMASH, TIC96-1038-C04001. Eduard Giménez and Juan A. Rodríguez-Aguilar enjoy the CIRIT doctoral scholarships 1998FI 0005 and FI-PG/96-8.490 respectively.

REFERENCES

[1] A. Chavez and Pattie Maes, 'Kasbah: An agent marketplace for buying and selling goods', in *First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96)*, pp. 75–90, (1996).

[2] D. Dubois, F. Esteva, P. Garcia, L. Godo, R. Lopez de Mantaras, and H. Prade, 'Fuzzy modelling of case-based reasoning and decision', in *Proceedings 2nd. Int. Conf. on Case Based Reasoning (ICCBR'97)*, eds., Leake and Plaza, pp. 599–611, (1997).

[3] D. Dubois and H. Prade, 'Possibility theory as a basis for qualitative decision theory', in *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, pp. 1924–1930, (1995).

[4] J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behaviour*, Princeton Univ. Press, Princeton, NJ, 1944.

[5] Juan A. Rodríguez-Aguilar, Francisco J. Martín, Pablo Noriega, Pere Garcia, and Carles Sierra, 'Competitive scenarios for heterogeneous trading agents', in *Proceedings of the Second International Conference on Autonomous Agents (AGENTS'98)*, (1998).

[6] Juan A. Rodríguez-Aguilar, Pablo Noriega, Carles Sierra, and Julian Padget, 'Fm96.5 a java-based electronic auction house', in *Second International Conference on The Practical Application of Intelligent Agents and Multi-Agent Technology(PAAM'97)*, (1997).

[7] M. Tsvetovatyy and M. Gini, 'Toward a virtual marketplace: Architectures and strategies', in *First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96)*, pp. 597–613, (1996).

[8] The FishMarket Project. http://www.iiia.csic.es/Projects/fishmarket.

[9] The Michigan AuctionBot Project. http://auction.eecs.umich.edu.