Bidding Strategies for Trading Agents in Auction-Based Tournaments

Pere Garcia, Eduard Giménez, Lluis Godo, and Juan A. Rodríguez-Aguilar

Artificial Intelligence Research Institute, IIIA
Spanish Council for Scientific Research, CSIC
08193 Bellaterra, Barcelona, Spain.
{pere,duard,godo,jar}@iiia.csic.es
http://www.iiia.csic.es

Abstract. Auction-based electronic commerce is an increasingly interesting domain for AI researchers. In this paper we present an attempt towards the construction of trading agents capable of competing in multi-agent auction markets by introducing both a formal and a more pragmatic 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 (bids) 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.

1 Introduction

Auctions are an attractive domain of interest for AI researchers in at least two areas of activity. On the one hand, we observe that the proliferation of on-line auctions in the Internet —such as Auctionline[1], Onsale[2], InterAUCTION[3], eBay[4] and many others— has established auctioning as a main-stream form of electronic commerce. Thus, agent-mediated auctions, and more generally agent-mediated institutions[18], appear as a convenient mechanism for automated trading, mainly due to the simplicity of their conventions for interaction when multi-party negotiations are involved, but also to the fact that on-line auctions may successfully reduce storage, delivery or clearing house costs in many markets. This popularity has spawned AI research and development in auction servers[28,

[1] http://www.auctionline.com
[2] http://www.onsale.com
[3] http://www.interauction.com
[4] http://www.eBay.com

23] as well as in trading agents and heuristics[9, 16]. On the other hand, auctions are not only employed in web-based trading, but also as one of the most prevalent coordination mechanisms for agent-mediated resource allocation problems (f.i. energy management[29], climate control[12], flow problems[27]).

¿From the point of view of multi-agent interactions, auction-based trading is deceivingly simple. Trading within an auction house demands from buyers merely to decide on an appropriate price on which to bid, and from sellers, essentially only to choose a moment when to submit their goods. But those decisions—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 on prices and participants' behaviour, etc. However, richness of information is not the only source of complexity in this domain. The actual conditions for deliberation are not only constantly changing and highly uncertain—new goods become available, buyers come and leave, prices keep on changing; no one really knows for sure what utility functions other agents have, nor what profits might be accrued—but on top of all that, deliberations are significantly time-bounded. Bidding times are constrained by the bidding protocol which in the case of DBP[5] auctions—like the traditional fish market[6]—proceeds at frenetic speeds.

Consequently, if a trading agent intends to behave aptly in this context, the agent's decision-making process may be quite elaborate. Clearly, the problem of choosing a successful bidding strategy by a trading agent in n-agents auction tournaments is clearly not deterministic and it will depend on many factors, in particular on the strategies themselves of the other competing agents. As long as the knowledge the agent will have about the other agents' strategies will be usually incomplete, our approach presented in this paper consists in looking at 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 over 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.

[5] *Downward bidding protocol*

[6] We will use the expression *fish market* to refer to the actual, real-world, human-based trading institution, and *FishMarket* to denote the artificial, formal, multi-agent counterpart. Hence, FM96.5 refers to a particular implementation of the *FishMarket* model of the fish market. Notice that we use the term *institution* in the sense proposed by North [19] as a "... *set of artificial constraints that articulate agent interactions*".

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) [17], 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 is possibilistic rather than probabilistic. Recently, Dubois and Prade have proposed in [7] 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[7] 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 work we build on this decision model to design bidding strategies for trading agents. After this introduction, Section 2 briefly describes the auction tournament environment. In Section 3, our theoretical decision model is described. Next section 4 is devoted to explain how this decision model can be applied to the design of bidding strategies while section 5 presents an heuristic-based approach for applying our model. Finally, we end up with some conclusions and an outline of our future work.

2 Auction-based Tournament Scenarios

In [23] we presented FM96.5, an electronic auction house based on the traditional *fish market* metaphor as an alternative approach to other proposals for electronic marketplace architectures(f.i. [25, 3]). 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

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

seller manager. The main 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 an *interagent*[15], 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. An interagent 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 test-bed, as described in [21, 22][8]. This test-bed 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 simplified[9] (formal) view of such tournament scenarios by introducing some of the elements composing them. 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*[10] and closes the round.

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

[8] The sof tware package of this test-bed can be donwloaded from the FM project web page[30] and has been fully reported in [20].

[9] The interested reader may refer to [22] for a more thorough discussion.

[10] The good is returned to its owner.

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 buyer's evaluation function calculated as the cumulative benefits.

It is worth noticing that a number of tournament scenarios of varying degrees of realism and complexity can be generated by instantiating the definition above[11].

3 Possibilistic-based Decision Theory

In this section we describe the possibilistic-based decision making model that we shall subsequently employ for designing competitive bidding strategies for trading agents.

We start by introducing the basics of Dubois and Prade's possibilistic decision model[7] (with some simplifications), and then we follow with some extensions that we propose in order to show how this decision model generalizes other decision models such as, f.i., Gilboa and Schmeidler's CBDT[10].

3.1 Background

First of all we introduce some notation and definitions. $X = \{x_1, \ldots, x_p\}$ will denote a finite set of consequences, $(V, \leq)$ a linear scale of uncertainty, with $\inf(V) = 0, \sup(V) = 1$. $Pi(X)$ 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 $\pi(x) = 1\}$. Finally, $(U, \leq)$ will denote a linear scale of preference (or utility), with $\sup(U) = 1$ and $\inf(U) = 0$, and $u : X \to U$ a utility function that assigns to each consequence x of X a preference level $u(x)$ of U. For the sake of simplicity here, we make the assumption that $U = V = [0, 1]$.

The working assumption of the decision model is that every decision $d \in D$ induces a possibility distribution $\pi : X \to V$ on the set X of consequences. Thus, ranking decisions amounts to ranking possibility distributions of $Pi(X)$. In such a framework, Dubois and Prade [7] propose the use of two kinds of qualitative utility functions to order possibility distributions. The basic underlying idea is

[11] [21] provide examples of several tournament instantiations.

based on the fact that a utility function $u : X \to U$ on the consequences can be regarded as specifying a fuzzy set of *preferred, good consequences*: the greater $u(x)$ is, the more preferred the consequence x is and the more x belongs to the (fuzzy) set of preferred consequences. On the other hand, a possibility distribution $\pi : X \to V$ specifies the fuzzy set of which consequences are plausible: the greater $\pi(x)$, the more plausible is the consequence x. Therefore, a conservative criterion is to look for those π's which, at some extent, make hardly plausible all the bad consequences, or in other words, all plausible consequences are good. On the contrary, an optimistic criterion that may be used to break ties is to look for those π's that, also to some extent, make plausible some of the good consequences.

For each utility function $u : X \to U$ the conservative and optimistic qualitative utilities used in the possibilistic decision model are respectively:

$$QU^-(\pi \mid u) = \min_{x \in X} \max(1 - \pi(x), u(x))$$
$$QU^+(\pi \mid u) = \max_{x \in X} \min(\pi(x), u(x)).$$

One can easily notice that $QU^-(\pi \mid u)$ and $QU^+(\pi \mid u)$ are nothing but the necessity and possibility degrees of the fuzzy set u w.r.t. the distribution π [1], or in other words, the Sugeno integrals of the utility function u with respect to the necessity and possibility measures induced by the distribution π. Moreover, when π denotes a crisp subset A (i.e. $\pi(x) = 1$ if $x \in A$, $\pi(x) = 0$ otherwise), $QU^-(\pi \mid u) = \min_{x \in A} u(x)$ and $QU^+(\pi \mid u) = \max_{x \in A} u(x)$, and hence, maximizing QU^- and QU^+ generalizes the well-known *maximin* and *maximax* decision criteria respectively. See [6] for an axiomatization of the preference relation induced by QU^-, QU^+, and other related utility functions.

3.2 Possible generalizations

It is well known in fuzzy set theory that the necessity and possibility measures account for a qualitative notion of fuzzy set inclusionship and intersection, respectively. Thus, in terms of fuzzy set operations, the decision criteria above using the QU^- and QU^+ functions can be read as the higher the degree of fuzzy set inclusionship of the π into u, the higher ranking of π according to the conservative criterion, while the higher the degree of fuzzy set intersection of π with u, the higher ranking of π according to the optimistic criterion.

Thus, besides those pure qualitative utilities, one can naturally think of introducing some other expressions of a more quantitative nature, but still accounting for a notion of inclusion and intersection. For instance, the most general way of defining the degree of intersection of π and u is:

$$dg(\pi \cap u) = \max_{x \in X} (\pi \cap u)(x),$$

where $(\pi \cap u)(x) = \pi(x) \otimes u(x)$, $\otimes$ being a t-norm[12] operation in [0, 1]. However, to define a degree of inclusion of π into u, there are at least two ways based on: (i)to what extent all elements of π are also elements of u; (ii) the proportion of elements of $\pi \cap u$ with respect to the elements of π. The former comes from a logical view while the latter comes from a conditioning view. They lead to the following expressions:

- $dg_l(\pi \subseteq u) = \min_{x \in X} \pi(x) \Rightarrow u(x)$,
  where $\Rightarrow$ is a many-valued implication[13] function,
- $dg_c(\pi \subseteq u) = \frac{\|\pi \cap u\|}{\|\pi\|}$,
  where $\| \, \|$ denotes fuzzy cardinality[14].

At this point, the following remarks are in order.

1. If both π and u define crisp subsets of consequences, then $dg_l(\pi \subseteq u)$ is either 1 or 0, while $dg_c(\pi \subseteq u)$ is nothing but the relative cardinality of π inside u, and for both, the degree is 1 only if $\pi \subseteq u$.
2. When $\otimes = \min$ and $\alpha \Rightarrow \beta = \max(1 - \alpha, \beta)$, we recover the qualitative utility functions: $dg_l(\pi \subseteq u) = QU^-(\pi \mid u)$ and $dg(\pi \cap u) = QU^+(\pi \mid u)$.
3. When $\otimes$ = *product*, $dg_c(\pi \subseteq u)$ is nothing but the expected value $E(u)$ of the utility function u w.r.t. to the unnormalized probability distribution $P(x) = \pi(x)$, or in other words, the weighted average of the $u(x)$ values according to the weights $\pi(x)$. When π comes from a similarity function, then $dg_c(\pi \subseteq u)$ can be closely related to Gilboa and Schmeidler's CBDT.

Finally, based on the notions of degree of inclusion and intersection defined above, we can consider the utility functions $\mathcal{U}^-_*(\pi \mid u) = dg_*(\pi \subseteq u)$, $* = l, c$, and $\mathcal{U}^+(\pi \mid u) = dg(\pi \cap u)$.

4 Possibilistic-based Design of Bidding Strategies

An agent's bidding strategy must decide on an appropriate price on which to bid for each good being auctioned during each round composing the tournament. Due to the nature of the domain faced by the agent, we must demand that such bidding strategy balances the agent's short-term benefits with its long-term benefits in order to succeed in long-run tournaments.

In what follows we make use of the possibilistic-based decision-making model described above as the key element to produce a competitive bidding strategy.

[12] A t-norm $\otimes$ is a binary operation (usually continuous) in [0, 1] which is non-decreasing, commutative, associative, and verifying $1 \otimes x = x$ and $0 \otimes x = 0$ for all $x \in [0, 1]$.

[13] An implication function $\Rightarrow$ is a binary operation in [0, 1] which is non-increasing in the first variable, non-decreasing in the second variable, and verifying at least $1 \Rightarrow x = x$ and $x \Rightarrow 1 = 1$ for all $x \in [0, 1]$.

[14] If A denotes a fuzzy subset of X with membership function μ_A then $\|A\| = \Sigma_{x \in X} \mu_A(x)$

4.1 The Decision Problem

For each round composing a tournament scenario, the decision problem for a trading agent consists in selecting a bid from the whole set of possible bids—from the starting price down to the reserve price.

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

$$s = (r, a, \tau, g, p_\alpha, p_{rsl}, \overline{\kappa}, \overline{E}, R)$$

characterizing round r of auction a such that τ is the type of the good g to be auctioned, p_α is its starting price, p_{rsl} is its resale price, $\overline{\kappa}$ is the vector of credits (κ_i is the 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 on the set of allowed bids our agent b_0 can submit. Given a new market situation s_0, we shall have $\mathcal{D} = \{bid(p) \mid p = p_\alpha - m.\Delta_{price}, m \in I\!N, p_{rsv} \leq p \leq \overline{\kappa}(b_0)\}$, where p_α and p_{rsv} are the starting and reserve prices in situation s_0, and $bid(p)$ means that the agent submits a bid at price p.

At each round, either the agent (b_0) wins, or buyer b_1 wins, ... , or buyer b_n wins by submitting bids at different prices. Therefore, the set X of outcomes (or consequences) is defined as the set $X = \{win(b_i, p) \mid i = 0, \ldots, n \;; p \in [p_{rsv} + \Delta_{price}, p_\alpha]\}$, where $x = win(b_i, p)$ means that buyer b_i wins the round by submitting a bid at 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 $\pi(x)$ 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.

Hereafter we shall assume that the agent keeps a memory of cases M storing the history of (past and the current) tournaments, whose cases are of the form $c = (s, b, p_s)$, where b is the buyer who won the round characterized by s (as defined above) by submitting a bid at price p_s.

4.2 Generating possibility distributions from cases

In order to obtain a possibility degree for each consequence in X, we observe the behaviour of each agent in previous similar situations. Then, the uncertainty on the behaviour of each agent in front of a new market situation is estimated, as a possibility degree, in terms of the similarity between the current situation and those market situations where the agent exhibited that behaviour.

Given the current market situation s_0, for each possible bid $p_d \in \mathcal{D}$, our agent has to evaluate the possibility of each buyer (including himself) winning the round, i.e. the possibility of each consequence $x \in X$. Let $x = win(b_i, p_0)$ be a consequence and (s, b_i, p) a case in M. We shall assume as a working principle that "the *more similar* is (s_0, p_0) to (s, p), the *more possible* b_i will be the winner in s_0" (a similar principle has been recently considered in a framework of fuzzy case-based reasoning[4]). If $\tilde{s}$ denotes the fuzzy set of situations similar to s, the above principle can be given the following semantics:

$$\pi_{s_0}(win(b_i, p_0)) \geq \mu_{\tilde{s}}(s_0) \otimes \mu_{\tilde{p}}(p_0)$$

where $\mu_{\tilde{s}} : Sit \rightarrow [0, 1]$ denotes the membership function of the fuzzy set $\tilde{s}$ and $\mu_{\tilde{p}} : Prices \rightarrow [0, 1]$ denotes the membership function of the fuzzy set $\tilde{p}$. They are defined as $\mu_{\tilde{s}}(s') = \mathcal{S}(s, s')$ and $\mu_{\tilde{p}}(p') = \mathcal{P}(p, p')$, where $\mathcal{S}$ and $\mathcal{P}$ are fuzzy relations on the set of situations and on the set of prices respectively, accounting for a notion of proximity or similarity.

Therefore, we can estimate the possibility degrees for each $b_i \neq b_0$ as:

$$\pi_{s_0}(win(b_i, p_0)) = \max_{\{(s,b_i,p) \in M | p \leq p_0\}} \mu_{\tilde{s}}(s_0) \otimes \mu_{\tilde{p}}(p_0)$$

for all $win(b_i, p_0) \in X$. Observe that this defnition ensures that the possibility of winning is non-increasing with respect to the value of bids (i.e. the lower the bid, the lesser the possibility of winning).

¿From these possibilities we can construct an initial fuzzy set $Bid^0_{b_i}$ of the possible winning bids of each participating buyer $b_i \neq b_0$ by defining its membership function as

$$\mu_{Bid^0_{b_i}}(p) = \pi_{s_0}(win(b_i, p))$$

for all p such that $win(b_i, p) \in X$. However this fuzzy set may be further modified by means of a set of fuzzy rules which attempt at modelling the rational behaviour of buyers in particular situations that may not be sufficiently described by the cases in the memory. For instance, we consider the following set of fuzzy rules:

if [$\overline{\kappa}(b_i)$ is *high*] and [R is *very_short*] and [$\overline{E}(b_i)$ is *low*]
then ΔBid_{b_i} is *very_positive*

if [$\overline{\kappa}(b_i)$ is *medium*] and [R is *very_short*] and [$\overline{E}(b_i)$ is *low*]
then ΔBid_{b_i} is *slightly_positive*

expressing heuristic rules describing expected changes in the strategy of a buyer when only a few rounds are left (R is *very_short*), and he lags behind in the ranking ($\overline{E}(b_i)$ is *low*). In these situations, depending on the agents' current credit ($\overline{\kappa}(b_i)$), the fuzzy rules above model an increase in the agresiveness of the buyer, at different degrees, by yielding the expected increases (ΔBid_{b_i}) in the agent's bid. In general, by applying a set of fuzzy rules of that type in the standard way, we obtain for each buyer a fuzzy set ΔBid_{b_i} representing the expected variation of the observed bidding strategy of each buyer.

¿From the combination of the initial fuzzy set of possible bids $Bid^0_{b_i}$ with the fuzzy set of expected variations ΔBid_{b_i} we obtain the final fuzzy set of possible bids

$$Bid^{\omega}_{b_i} = Bid^0_{b_i} \oplus \Delta Bid_{b_i}$$

where $\oplus$ denotes fuzzy addition, i.e.

$$\mu_{Bid^{\omega}_{b_i}}(p) = \max\{\min\{\mu_{Bid^{0}_{b_i}}(p_1), \mu_{\Delta Bid_{b_i}}(p_2)\} \mid p = p_1 + p_2\}.$$

Then, we make use of the fuzzy set Bid^{ω} to reassign possibilities to each consequence for each $b_i \neq b_0$

$$\pi_{s_0,p_d}(win(b_i,p)) = \begin{cases} \mu_{Bid^{\omega}(b_i)}(p), & \text{if } 0 < p < p_d \\ 0, & otherwise \end{cases}$$

Finally, to estimate the possibility of our agent winning with a bid at price p_d we look into the memory M for those cases such that the sale price was not greater than p_d. Let $M_{p_d} = \{(s, b_i, p) \in M \mid p < p_d, b_i \neq b_0\}$. Then

$$\pi_{s_0,p_d}(win(b_0,p)) = \begin{cases} \max\limits_{(s,b_i,p') \in M_{p_d}} \mu_{Bid^{\omega}(b_i)}(p'), & \text{if } p = p_d \\ 0, & otherwise \end{cases}$$

These are the possibility values to be utilized when applying our decision model.

4.3 Assessing utilities

Given a new market situation s_0, for each consequence $x = win(b_i, p)$ our agent b_0 must assess the utility value $u(win(b_i, p))$ at the fact that buyer b_i wins the round by submitting a bid at price p. Several modelling options could be considered here. As a matter of example we propose here a particular utility function that 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 the 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 $k = (\max_{j \neq i} \overline{E}(b_j)) - \overline{E}(b_i)$ and $t = (R-1)/(\max(\overline{\kappa}(b_i) - p, 1) \cdot (p_{rsl} - p))$, being p_{rsl} the resale price. We assume that $p_{rsl} - p \geq 0$, and $\overline{\kappa}(b_i) - p \geq 0$, i.e., buyers only take into consideration bids that can improve their score whenever they have enough credit to submit them. In the above definition of f, the factor $(max_{j \neq i} \overline{E}(b_j)) - \overline{E}(b_i)$ accounts for the position of buyer b_i with respect to the other buyers in the ranking of scores, the factor $p_{rsl} - p$ accounts for the benefit the agent would make if he wins the round, and the factor $\frac{R-1}{max(\overline{\kappa}(b_i)-p,1)}$ estimates the cost of winning the round. Then, based on the scoring function f, we propose the following utility function:

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

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

To summarize, given a new market situation s_0, the decision process follows the following steps:

(i) for each decision $d = bid(p_d)$, where $p_d \in \mathcal{D}$,
   (a) for each consequence $x = win(b_i, p) \in X$ we calculate:
   - the possibility $\pi_{s_0,d}(x)$
   - the utility $u(x)$
   (b) the global utility assessed to each decision d will be calculated from either QU^-, QU^+, $\mathcal{U}^-$ or $\mathcal{U}^+$ by combining possibilities with utilities.

(ii) Our agent b_0 will choose one of the most preferred decision(one decision valued most by the global utility function).

5 An Heuristic Approach for the Development of Competitive Bidding Strategies

For each round, the resulting strategy performs a hybrid, two-fold decision making process that involves the usage of global(market-centered) probabilistic information in a first decision step, and individual(rival-centered) possibilistic information in a second, refining decision step.

The outright use of the possibilistic 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 market. Then, the second step consists in 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 [5]. This amounts to consider, for each case[15] $(s, p) \in M$, 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 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 4.2, since the criteria used to define how similar situations are 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 [8]) for details about the semantics of fuzzy gradual rules):

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

[15] 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.

Here I is residuated many-valued implication[16], 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 $\hat{B}_\alpha = \cap_{(s,p) \in M} \{p' \mid T^*(p, p') \geq S^*(s, s_0) \otimes \alpha\}$, 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 algorithm in figure 1 summarizes the process of selecting a bid out of $\hat{B}_\alpha$ using a utility function U.

6 Conclusions and Future Work

We have described a possibilistic-based decision method that attemps at modelling buyer agents' behaviour in electronic auction tournaments. Interestingly, competitions seem to be in vogue in the AI community as suggested by the many emerging initiatives. *Robocup*[13] is attempting to encourage both AI researchers and robotics researchers to make their systems play soccer, autonomous mobile robots try to show their skills in office navigation and in cleaning up the tennis court in the *AAAI Mobile Robot Competition*[14], and even automated theorem proving systems participate in competitions [24]. But surely our proposal is closer to the *Double auction* tournaments held by the Santa Fe Institute[2] where the contenders competed for developing optimized trading strategies. However, the main concern of our proposal consists in providing a method for performing multi-agent reasoning under uncertainty based on the modelling of the other agents' behaviour likewise [26], where the recursive modelling method [11] was used for constructing agents capable of predicting the other agents' behaviour in Double auction markets.

As to our future work, firstly this shall focus on the empirical evaluation of our proposal. Secondly, we will head towards the construction of actual agents capable of trading in actual auction markets under the rules of any bidding protocol.

[16] 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 $x \otimes 0 = 0$.

Function **Bid_Selection** $(M, \hat{B}_\alpha, s_0, U)$
$\underline{\forall}$ candidate bid $p_d \in \hat{B}_\alpha$
$\underline{\forall}$ buyer $\mathbf{b} \in \mathcal{B}$ such that $b \neq b_0$
Retrieve all cases $c = (s, \mathbf{b}, p_s) \in M$
$\forall$ price p
Consequence $x := win(\mathbf{b}, p)$
if $p < p_d$
then $\pi(win(\mathbf{b}, p)) = 0$
else $\pi(win(\mathbf{b}, p)) = \pi_{s_0, p_d}(win(b, p))$
$u(win(\mathbf{b}, p)) = r(-f(\mathbf{b}, s_0, p))$
end for
end for
$b = b_0$
Retrieve all cases $c = (s, b_0, p_s) \in M$
$\forall$ price p
Consequence $x := win(\mathbf{b_0}, p)$
if $p \neq p_d$
then $\pi(win(\mathbf{b_0}, p)) = 0$
else $\pi(win(\mathbf{b_0}, p)) = \pi_{s_0, p_d}(win(b_0, p))$
$u(win(\mathbf{b_0}, p) = r(f(\mathbf{b_0}, s_0, p))$
end for
Calculate $U(bid(\mathbf{b_0}, p))$
end for
return $p = \arg\max\{U(bid(\mathbf{b_0}, p_d)) \mid p_d \in \hat{B}\}$
end function

Fig. 1. Bid Selection Procedure

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. Zadeh L. A. Fuzzy sets as a basis for the theory of possibility. *Fuzzy Sets and Systems*, (1):3–28, 1978.
2. M. Andrews and R. Prager. *Genetic Programming for the Acquisition of Double Auction Market Strategies*, pages 355–368. The MIT Press, 1994.
3. 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)*, pages 75–90, 1996.
4. D. Dubois, F. Esteva, P. Garcia, L. Godo, R. Lopez de Mantaras, and H. Prade. Fuzzy modelling of case-based reasoning and decision. In Leake and Plaza, editors, *Proceedings 2nd. Int. Conf. on Case Based Reasoning (ICCBR'97)*, pages 599–611, 1997.

5. D. Dubois, F. Esteva, P. Garcia, and H. Prade. A logical approach to interpolation based on similarity relations. *Journal of Approximate Reasoning*, 17(1):1–36, 1997.
6. D. Dubois, Lluis Godo, Henri Prade, and Adriana Zapico. Making decision in a qualitative setting: from decision under uncertainty to case-based decision. In *Proceedings of the 6th.. Int. Conf. on Principles of Knowledge Representation and Reasoning(KR'98)*, 1998.
7. 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)*, pages 1924–1930, 1995.
8. D. Dubois and H. Prade. What are fuzzy rules and how to use them. *Fuzzy Sets and Systems*, 84:169–185, 1996.
9. Pere Garcia, Eduard Giménez, Lluís Godo, and Juan A. Rodríguez-Aguilar. Possibilistic-based design of bidding strategies in electronic auctions. In *The 13th biennial European Conference on Artificial Intelligence (ECAI-98)*, 1998.
10. I. Gilboa and D. Schmeidler. Case-based theory. *The Quarterly Journal of Economics*, 110:607–639, 1995.
11. Piotr Gmytrasiewicz and Edmund H. Durfee. A rigorous, operational formalization of recursive modeling. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 125–132, 1995.
12. B. A. Huberman and S. Clearwater. A multi-agent system for controlling builging environments. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 171–176. AAAI Press, June 1995.
13. Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, and Eiichi Osawa. Robocup: The robot world cup initiative. In *First International Conference on Autonomous Agents*, 1997.
14. David Kortenkamp, Illah Nourbakhsh, and David Hinkle. The 1996 AAAI Mobile Robot Competition and Exhibition. *AI Mag.*, 18(1):25–32, 1997.
15. Francisco J. Martín, Enric Plaza, and Juan Antonio Rodríguez-Aguilar. An infrastructure for agent-based systems: An interagent approach. *International Journal of Intelligent Systems*, 1998.
16. Noyda Matos, Carles Sierra, and Nick R. Jennings. Determining successful negotiation strategies: An evolutionary approach. In *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*, 1998.
17. J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behaviour*. Princeton Univ. Press, Princeton, NJ, 1944.
18. Pablo Noriega. *Agent-Mediated Auctions: The Fishmarket Metaphor*. PhD thesis, Universitat Autonoma de Barcelona, 1997. Also to appear in IIIA mongraphy series.
19. D. North. *Institutions, Institutional Change and Economics Perfomance*. Cambridge U. P., 1990.
20. Juan A. Rodríguez-Aguilar, Francisco J. Martín, Francisco J. Giménez, and David Gutiérrez. Fm0.9beta users guide. Technical report, Institut d'Investigació en Intel.ligència Artificial. Technical Report, IIIA-RR98-32, 1998.
21. 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)*, pages 293–300, 1998.
22. Juan A. Rodríguez-Aguilar, Francisco J. Martín, Pablo Noriega, Pere Garcia, and Carles Sierra. Towards a test-bed for trading agents in electronic auction markets. *AI Communications*, 11(1):5–19, 1998.

23. 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)*, pages 207–224, 1997.
24. Christian B. Suttner and Geoff Sutcliffe. *ATP System Competition*, volume 1104 of *Lecture Notes in Artificial Intelligence*, pages 146–160. Springer Verlag, 1996.
25. 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)*, pages 597–613, 1996.
26. José M. Vidal and Edmund H. Durfee. Building agent models in economic societies of agents. In *Workshop on Agent Modelling (AAAI-96)*, 1996.
27. Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, (1):1–23, 1993.
28. Peter R. Wurman, , Michael P. Wellman, and William E. Walsh. The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents. In *Second International Conference on Autonomous Agents (AGENTS'98)*, 1998.
29. Fredrik Ygge and Hans Akkermans. Making a case for multi-agent systems. In Magnus Boman and Walter Van de Velde, editors, *Advances in Case-Based Reasoning*, number 1237 in Lecture Notes in Artificial Intelligence, pages 156–176. Springer-Verlag, 1997.
30. The FishMarket Project. http://www.iiia.csic.es/Projects/fishmarket.