Designing Bidding Strategies for Trading Agents in Electronic Auctions

Eduard Giménez-Funes Lluís Godo Juan A. Rodríguez-Aguilar
Pere Garcia-Calvés
Artificial Intelligence Research Institute, IIIA
08193 Bellaterra, Barcelona, Spain
{duard,godo,jar,pere}@iiia.csic.es

Abstract

Auction-based electronic commerce is an increasingly interesting domain for developing trading agents. In this paper we present our first contributions towards the construction of such agents by introducing both a formal and a more pragmatical approach for the design of bidding strategies that provide buyer agents with useful heuristic guidelines to participate in auction-based tournaments. On the one hand, our formal view relies on possibilistic-based decision theory as the means of handling possibilistic uncertainty on the consequences of actions due to the lack of knowledge about the other agents' behaviour. On the other hand, for practical reasons we also 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) probabilistic information in a first decision step which is subsequently refined by a second decision step based on the individual (rival-centered) possibilistic information induced from the memory of cases composing the history of tournaments. In this way, the resulting bidding strategy balances the agent's short-term benefits, related to the probabilistic information, with its long-term benefits, related to the possibilistic information.

1. Introduction

Online auctions such as Auctionline, Onsale, InterAUCTION, eBay and many others have proliferated over the Internet as well-established mechanisms for multi-party negotiation. As a matter of fact, 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, the implication of many traders—instead of a conventional sale's single buyer and seller—, and the fact that online auctions do significantly reduce costs. As a major benefit of dynamically negotiating a price through auctions, the task of determining the value of a good is transferred from the merchant to the market, leading to a fair allocation of limited resources (to those who value them most). This is the reason why auctions are not only employed in online retail, but also as the key element for building solutions to resource allocation problems(f.i. energy management[20], climate control[8], flow problems[19]). Hence, auctions must be regarded as an attractive domain for developing agents.

Nonetheless, designing, building, and tuning trading agents before letting them loose in wildly competitive scenarios like electronic auctions inhabited by (both human and software) expert traders happens to be an arduous task. This fact motivated the construction of FM97.6, a test-bed for electronic auctions [14] that intends to provide support to agent developers in such a challenge. 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.

Trading within an auction house demands from buyers 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, other 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. Consequently, if a trading agent intends to behave aptly in this context, the agent's decision-making process may be quite elaborate.

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 Theory problem, the trading agent has to choose a decision (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 possible states or situations Sit and a (finite) set X of possible consequences or outcomes of the decisions, a decision d is represented by a function $d : S \rightarrow X$ assigning to each situation s the consequence $d(s)$ of having taken the decision d at the state s. On the other hand, consequences are ranked by a utility function $u : X \rightarrow I\!R$ modelling the decision maker's preferences among them. In a decision process, uncertainty may be involved in knowing either what the real situation is or what the precise consequences of decisions are. Classical approaches to decision making under uncertainty assume that uncertainty is represented by probability distributions. In the first case it is assumed that a probability distribution P on the set of situations is known. Then, the utility of a decision $d : Sit \rightarrow X$ is measured by the expected value of u w.r.t. to P:

$$U(d) = \Sigma_{s \in Sit} P(s) u(d(s)).$$

In the case when we know the real situation but the consequences of decisions are not precisely known (for each decision d we only know the probability $P_d(x)$ of each consequence $x \in X$), then the utility of a decision d is evaluated in an analogous way as the expected value of u w.r.t. to P_d:

$$U(d) = \Sigma_{x \in X} P_d(x) u(x).$$

This kind of approaches correspond to the well-known Expected Utility Theory (EUT) [12, 16], but they present some problems and paradoxes, basically related with inferring the probabilities. Indeed, EUT is well suited when the space S of the possible states is well-known, it is easy to figure out which is the outcome of each decision at each state, and of course when the probability distributions over states or outcomes are also well-known. However, 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 successful bids corresponding to previous tournaments. This is precisely the kind of decision problems addressed by the so-called Case-based Decision Theory (CBDT) proposed by Gilboa and Schmeidler in [6], where cases are viewed as instances of decision making and where the decision maker only knows cases which have been previously experienced. In that case, given a memory M of precedent decision problem instances, represented by triples $(situation, decision, outcome)$, and a real-valued similarity function S among situations, they propose in a new situation s_0 to choose the decision which maximises the following expression

$$U(d) = \Sigma_{(s,d,x) \in M} S(s_0, s) u(x).$$

In [3, 4] another case-based view of decision making is proposed and shown to be in close connection with a possibilistic decision model earlier introduced by Dubois and Prade in [5]. Indeed, the process of looking at the winning behaviour of other trading agents in previous similar situations and adopting a similar behaviour for our agent induces an uncertainty which is of possibilistic nature rather than probabilistic: the more similar are the cases, the more plausible is to assume the outcome of a given decision in one case for the other one.

In this paper we adapt this latter decision model to design bidding strategies for trading agents. Next section sketches out the auction tournament environment FM97.6, and presents how tournament scenarios can be formally defined. In Section 3, the theoretical underpinnings of the decision model are described, while Section 4 is devoted to explain how this decision model can be employed in the design of bidding strategies for tournament scenarios. Finally, we end up with a discussion about related work and an outline of our future work.

2. A Test-bed for Auction Tournaments

Following [13], the fish market can be described as a place where several *scenes* run simultaneously, at different places, but with some causal continuity. 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 according to the *downward bidding protocol* whose dynamics is described as follows:

[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* and closes the round.

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

However, before those boxes of fish may be sold, fishermen have to deliver the fish to the fish market, at the *sellers' registration scene*, and buyers need to register for the market, at the *buyers' registration scene*. Likewise, once a box of fish is sold, the buyer should take it away by passing through a *buyers' settlements scene*, while sellers may collect their payments at the *sellers' settlements scene* once their lot has been sold.

In [15, 21, 13] we present an electronic auction house based on the traditional fish market metaphor. In a highly mimetic way, the workings of FM96.5 also involve the concurrency of several scenes governed by the market intermediaries identified in *FishMarket*. Therefore, seller agents register their goods with a seller admitter agent, and can get their earnings (from a seller manager) once the auctioneer has sold these goods in the auction room. Buyers, on the other hand, register with a buyer admitter, and bid for goods which they pay through a credit line that is set up and updated with a seller manager. Buyer and seller agents can trade goods as long as they comply with the *FishMarket institutional* conventions. Those conventions that affect buyers and sellers have been coded into what we call a *market interagent*[11] 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 market interagent gives a permanent identity to the trader and enforces a *conversation protocol* that establishes what illocutions can be uttered by whom and when —and consequently what their language and content, sequencing and effects may be[1].

In order to obtain an auction tournament environment, more functionality has been added to FM96.5 to turn it into a test-bed, FM97.6. The resulting test-bed has the following salient characteristics[2]:

- It is *domain-specific* in the sense that it models and simulates an *electronic auction house*.
- It is *realistic*, since it follows the actual conventions of a complex real-world institution, the traditional fish market.
- The use of market interagents makes FM97.6 *architecturally–neutral* since no particular agent architecture (or language) is assumed or provided.
- FM97.6 allows for very flexible *scenario generation* to enable designers to produce systematic experimentation. FM97.6 allows for the specification, and subsequent activation, of a large variety of market scenarios: from very simple artificial scenarios to complex realistic scenarios.
- Explicit parameter-fixing and participant-registration modes are involved in the scenario generation facility, to allow for the *repeatability* of experiments.
- A trace tool keeps track of all illocutions and transactions that take place during an auction. Hence, a whole auction can be audited and re-enacted step-by-step, and the evolving performance of all the agents involved in a tournament can be traced, evaluated, and analyzed.

Summarizing, the resulting environment, FM97.6, constitutes a test-bed where a very rich variety of experimental conditions, tournament scenarios, can be explored systematically and repeatedly, and analyzed and reported with lucid detail if needed.

Each one of these tournament scenarios 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 this section we identify the elements composing tournament scenarios by introducing our formal notion of *tournament descriptor*.

The type of tournaments that we do consider follow the downward bidding protocol described in the section above. Though the protocol thus defined is vague, notice, however, that a finite set of parameters that control the dynamics of the bidding process are implicit in this protocol definition. Hence, we formally define a **Downward Bidding Protocol Dynamics Descriptor** $\mathcal{D}_{DBP}$ as the tuple $\langle \Delta_{price}, \Delta_{offers}, \Delta_{rounds}, \Sigma_{coll}, \Pi_{sanction}, \Pi_{rebid} \rangle$ such that

- $\Delta_{price} \in I\!N$ (price step). Decrement of price between two consecutive quotations uttered by the auctioneer.
- $\Delta_{offers} \in I\!N$ (minimum time between offers). Delay between consecutive price quotations.

[1] In [15] we used the term *nomadic agent interface*; in [13, Chpt.10] the notion of *institutor* is defined and discussed.

[2] Refer to [14] for a more thorough discussion.

- $\Delta_{rounds} \in I\!N$ (minimum time between rounds). Delay between consecutive rounds belonging to the same auction.

- $\Sigma_{coll} \in I\!N$ (maximum number of successive collisions). This parameter prevents the algorithm from entering an infinite loop as explained above.

- $\Pi_{sanction} \in I\!R$ (sanction factor). This coefficient is utilized by the buyers' manager to calculate the amount of the fine to be imposed on buyers submitting unsupported bids.

- $\Pi_{rebid} \in I\!R$ (price increment). This value determines how the new offer is calculated by the auctioneer from the current offer when either a collision, or an unsupported bid occur.

Note that the identified parameters impose significant constraints on the trading environment. For instance, Δ_{offers} and Δ_{rounds} affect the agents' time-boundedness, and consequently the degree of situatedness viable for bidding strategies.

Next we introduce the notion of tournament descriptor as an attempt to encompass all the information characterizing tournament scenarios. Thus, we define a **Tournament Descriptor** $\mathcal{T}$ as the tuple $\langle n, \Delta_{auctions}, D, P, B, S, F, E\rangle$ such that:

- n is the number of auctions to take place during a tournament.

- $\Delta_{auctions}$ is the time between consecutive auctions.

- D is a finite set of bidding protocols' dynamics descriptors.

- P is a finite family of communication protocols that a buyer agent must employ to interact with its *interagent* indexed by different bidding protocol types (f.i. $P = \{P_{DBP}, P_{English}, \ldots\}$).

- $B = \{b_1, \ldots, b_p\}$ is a finite set of identifiers corresponding to all participating buyers.

- $S = \{s_1, \ldots, s_q\}$ is a finite set of identifiers corresponding to all participating sellers.

- $F = [\mathcal{F}^1, \ldots, \mathcal{F}^n]$ is a sequence of n descriptors. Each $\mathcal{F}^i$ specifies the way auction $\mathcal{A}^i$ is dynamically generated.

- $E = \langle E_b, E_s\rangle$ is a pair of winner evaluation function that permit to calculate respectively the score of buyers and sellers.

Observe that a multitude of experimental tournament scenarios of varying degrees of realism and complexity can be generated by the tournament designer when instantiating the definition of tournament descriptor[3]. The information within the *tournament descriptor* must be conveyed to the buyers participating in tournaments so that they know the features of the competitive scenario they are immersed in.

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[5] (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.

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 $\pi_d : X \to V$ on the set X of consequences. Thus, ranking decisions amounts to ranking possibility distributions of $Pi(X)$. distribution. In such a framework, Dubois and Prade [5] propose the use of two kinds of qualitative utility functions to order possibility distributions. The basic underlying idea is 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 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 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

[3][14] provides an example of a tournament instantiation

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 [4] 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 the π_d 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[4] 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[5] function,
- $dg_c(\pi \subseteq u) = \frac{\|\pi \cap u\|}{\|\pi\|}$,
  where $\| \, \|$ denotes fuzzy cardinality[6].

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. Case-based Decision Model for Designing 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. 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.

[4] 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 $x \in [0, 1]$.

[5] 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 $x \in [0, 1]$.

[6] 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

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

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 ($\kappa_i = \mathcal{C}(b_i)$), $\overline{E}$ is the vector of scores ($E_i = E_b(b_i)$), 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 submit. Given a new market situation s_0, we shall have $\mathcal{D} = \{bid(p) \mid p = p_\alpha - m.\Delta_{price}, m \in \mathbb{N}, p_{rsv} \le p \le \mathcal{C}(b_0)\}$, 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 $X = \{win(b_i, p) \mid i = 0, \dots, n\,; p \in [p_{rsv} + \Delta_{price}, p_\alpha]\}$, where $x = win(b_i, p)$ means that buyer b_i wins the round by submitting a bid at price p.

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 $c_{a,r} = (s_{a,r}, b, p_s)$, where *b* is the buyer who won the round characterized by $s_{a,r}$ (as defined above) by submitting a bid at price p_s.

Finally, we must recall from the decision model introduced in the previous section that given a new market situation s_0, the agent has to assess, for each possible decision (bid) $d \in \mathcal{D}$, the possibility and utility values $\pi_d(x)$ and $u(x)$, $\forall x \in X$, in order to be able to calculate a global utility for each d (using either QU^-, QU^+, $\mathcal{U}^-$, or $\mathcal{U}^+$). The way of generating possibilities and assessing utilities is presented along the next subsections.

4.2. Reducing the Search Space

Evidently, deploying the possibilistic-based decision mechanism from the whole set of possible decisions(bids) $\mathcal{D}$ appears to be prohibitively expensive. Instead, we reduce the decision set by considering a subset composed of those decisions(bids) maximizing the agent's short-term expected benefit for the current round, the so-called *set of candidate bids*. This pre-processing of $\mathcal{D}$ will ideally help the agent's deliberation process to constrain to time and resource-boundedness.

In order to obtain a set of candidate bids for a given round r of auction a characterized by a feature vector s, we firstly infer a probability distribution on the sale price P_s from the past history of the tournament. Secondly, we utilize such distribution to obtain the price $\tilde{p}$ which maximizes the agent's short-term expected benefit for the current round given by the following expression

$$expected\_benefit_{a,r}(p) = (p_{rsl} - p) Prob(P_s \le p)$$

where p_{rsl} is the resale price of the good to be auctioned, P_s follows a normal distribution $\mathcal{N}(\hat{p}, \sigma)$ such that the expected sale price $\hat{p}$ and the standard deviation σ are estimated by regression analysis on the memory of cases.

Finally, we will construct the set of candidate bids *Bids* by selecting the δ closest bids to $\tilde{p}$:

$$Bids = \{\tilde{p} \pm \delta' \Delta_{price} | \delta' = 0, 1, \dots, \frac{\delta - 1}{2}\}.$$

From this set, we shall redefine the decision set as $\mathcal{D} = \{bid(p) \mid p \in Bids\}$.

4.3. Generation of Possibility Distributions

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 $p_d \in Bids$, 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 bee recently considered in a framework of fuzzy case-based reasoning[3]). If $\tilde{s}$ denotes the fuzzy set of situations similar to s, the above principle can be given the following semantics:

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

where $\mu_{\tilde{s}} : Sit \to [0,1]$ denotes the membership function of the fuzzy set $\tilde{s}$ and $\mu_{\tilde{p}} : Prices \to [0,1]$ 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:

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

for all $win(b_i, p_0) \in X$. From these possibilities we can construct an initial fuzzy set $Bid^0_{b_i}(p)$ of the possible winning bids of each participating buyer $b_i \neq b_0$ as

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

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 [$\mathcal{C}(b_i)$ is *high*] and [R is *very_short*] and [$E_b(b_i)$ is *low*]
then ΔBid_{b_i} is *very_positive*

if [$\mathcal{C}(b_i)$ is *medium*] and [R is *very_short*] and [$E_b(b_i)$ 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 ($E_b(b_i)$ is low). In these situations, depending on the agents' current credit ($\mathcal{C}(b_i)$), 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 $Bid^0_{b_i}(p)$ 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.

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

Finally, we make use of the fuzzy set $Bid^{\omega}(b_i)$ to reassign possibilities to each consequence for each $b_i \neq b_0$

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

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

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

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

4.4. 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 at the fact that buyer b_i wins the round by submitting a bid at price p, $u(win(b_i, p))$. In what follows we propose a utility function for constructing 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.

For this purpose, we consider the following function:

$$f(b_i, s_0, p) = \begin{cases} k.\frac{R^0-1}{\max(\kappa^0_i - p, 1)}, & \text{if } k \leq 0 \\ k.\frac{\max(\kappa^0_i - p, 1)}{R^0-1}, & \text{otherwise} \end{cases}$$

where $k = ((\max_{j\neq i} E^0_j) - E^0_i)/(p^0_{rsl} - p)$, being p^0_{rsl} the resale price, and E the evaluation function for buyers. We assume that $p^0_{rsl} - p \geq 0$, and $\kappa^0_i - p \geq 0$, i.e. buyers only consider bids that can improve their score, and they have enough credit to submit the bid at price p. In f the factors $(max_{j\neq i} E^0_j) - E^0_i$, $p^0_{rsl} - p$, and $\frac{R^0-1}{max(\kappa^0_i - p, 1)}$ stand for the position of buyer b_i with respect to the other buyers in the ranking of scores, the net profit and the estimated cost of winning the round, respectively.

And from f, we define the utility function

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

where r is a linear scaling function which makes u fall into [0,1]. Notice that f is decreasing with respect to the number of rounds left R^0, but increasing with respect to the credit κ^0_i of buyer b_i. This means that the less rounds are left and the more money the buyer has got, the more the buyer will prefer to bid. On the other hand, when the buyer is in the lead the utility of bidding is valued negatively ($u \in [0, \frac{1}{2}]$), otherwise—the buyer is behind the leader–the utility of bidding is valued positively($u \in (\frac{1}{2}, 1]$).

At this point, we do have all the ingredients for applying the decision model proposed. To summarize, given a new market situation s_0, for each decision $d = bid(p_d)$, $p_d \in Bids$, we calculate the possibility $\pi_{s_0,d}(x)$ and utility $u(x)$ of each consequence $x = win(b_i, p) \in X$. Then, according to our decision model 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. Our agent b_0 will choose the most preferred decision(the decision valued most by the global utility function).

5. Related and Future Work

Interestingly, competitions seem to be in vogue in the AI community as suggested by the many emerging initiatives.

Robocup[9] 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*[10], and even automated theorem proving systems participate in competitions [17]. 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 [18], where the recursive modelling method [7] was used for constructing agents capable of predicting the other agents' behaviour in Double auction markets.

At present, a proof-of-concept implementation of our proposal is undergoing empirical evaluation. We are basically analyzing which utility and similarity functions yield good performances. In general, conservative utilities $\mathcal{U}_l^-$ lead to a preferring higher bids than $\mathcal{U}_c^-$. As to our future work, firstly our research will head towards the construction of actual agents capable of trading in actual auction markets under the rules of any auction protocol. Secondly, in parallel, FM97.6 will be made to evolve to host other (even more flexible) forms of price-fixing mechanisms (English auction, Double auction, discounting, open negotiation, etc.), and will be equipped with a trading-agent shell to help agent designers construct their agents.

6. 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] Z. 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] D. Dubois, F. Esteva, P. Garcia, L. Godo, R. L. 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.

[4] D. Dubois, L. Godo, H. Prade, and A. 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.

[5] 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.

[6] I. Gilboa and D. Schmeidler. Case-based theory. *The Quarterly Journal of Economics*, pages 607–639, 1995.

[7] P. Gmytrasiewicz and E. H. Durfee. A rigorous, operational formalization of recursive modeling. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 125–132, 1995.

[8] 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.

[9] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, and E. Osawa. Robocup: The robot world cup initiative. In *First International Conference on Autonomous Agents*, 1997.

[10] D. Kortenkamp, I. Nourbakhsh, and D. Hinkle. The 1996 AAAI Mobile Robot Competition and Exhibition. *AI Mag.*, 18(1):25–32, 1997.

[11] F. J. Martín, E. Plaza, J. A. Rodríguez-Aguilar, and J. Sabater. Jim a java interagent for multi-agent systems. In *Proceedings of the AAAI-98 Workshop on Software Tools for Developing Agents*, 1998.

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

[13] P. Noriega. *Agent-Mediated Auctions: The Fishmarket Metaphor*. PhD thesis, Universitat Autonoma de Barcelona, 1997. Also to appear in IIIA mongraphy series.

[14] J. A. Rodríguez-Aguilar, F. J. Martín, P. Noriega, P. Garcia, and C. Sierra. Competitive scenarios for heterogeneous trading agents. In *Proceedings of the Second International Conference on Autonomous Agents (AGENTS'98)*, 1998.

[15] J. A. Rodríguez-Aguilar, P. Noriega, C. Sierra, and J. 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.

[16] L. J. Savage. *The Foundations of Statistics*. Dover, 1972.

[17] C. B. Suttner and G. Sutcliffe. *ATP System Competition*, volume 1104 of *Lecture Notes in Artificial Intelligence*, pages 146–160. Springer Verlag, 1996.

[18] J. M. Vidal and E. H. Durfee. Building agent models in economic societies of agents. In *Workshop on Agent Modelling (AAAI-96)*, 1996.

[19] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, (1):1–23, 1993.

[20] F. Ygge and H. Akkermans. Power load management as a computational market. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96)*, 1996.

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