2

Constraint Based Automated Multi-attribute Negotiations

Miguel A. López-Carmona, Iván Marsá-Maestre and Juan R. Velasco
Universidad de Alcalá
Spain

1. Introduction

Multi-attribute (Multi-issue) negotiation protocols have been studied widely, and represent an important challenge in the multiagent systems community (Lai et al., 2004). Therefore, a lot of automated negotiation models and protocols have been developed, and manifold negotiation challenges have been already addressed. Most research in automated negotiation to date has focused on the competitive aspect (Vo et al., 2007). On the other hand, work by Dispute Resolution theorists in the social sciences has also focussed substantially on how to achieve negotiated agreements that are of a high value to all parties (Fischer & Ury, 1981). This approach is known as *Integrative* or *Interest-based negotiation*, and it has been recognised as the more successful approach to the negotiation problem. Example scenarios where such cases may arise are: business process management involving agents within the same organization, e-commerce negotiations where the seller is interested in having a satisfied buyer (e.g. long-term commercial relationships), or e-commerce scenarios where risk averse agents avoid the conflict in the negotiation processes. In the context of purchase negotiation scenarios, it is clear that every negotiation partner tries to maximize his preferences. However, when an agent aims at optimizing his own benefit with no regard for the others', it has been shown that negotiators more often than not reach inefficient compromises. Conflict theorists Lax and Sebenius (Lax & Sebenius, 1992) argue that negotiation necessarily includes both cooperative and competitive elements, and that these elements exist in tension. Therefore, he refers to the problem of deciding whether to pursue a cooperative or a competitive strategy at a particular time during a negotiation as the Negotiator's Dilemma. However, it is not always possible to separate the integrative bargaining process, i.e. when agents use cooperative strategies to search for joint gains, from the distributive bargaining process, i.e. when agents use competitive strategies in order to 'claim value'. The main problem is that distributive and integrative processes interplay with each other making information manipulation becomes part of the integrative bargaining process. *Integrative negotiation* contrasts with *distributive bargaining* in which the parties are trying to distribute a fixed resource, and where if an agent wins another agent looses. Distributive negotiation predicts that one party can only gain at the other party's expense. The key characteristics that distinguish integrative negotiations from distributive ones are: creation of value; focus on interests and not positions; openness and exchange of relevant

Source: Multiagent Systems, Book edited by: Salman Ahmed and Mohd Noh Karsiti,
ISBN 978-3-902613-51-6, pp. 426, February 2009, I-Tech, Vienna, Austria

information, and even learning; and problem restructuring. In order to achieve integrative approaches, literature of automated negotiation proposes a number of techniques such as *multi-attribute utility theory*, *distributed constraint satisfaction*, and *cojoint analysis*. A common aspect in all these techniques and in integrative negotiation approaches in general is that a multi-attribute negotiation scenario is required. Attributes are the characteristics of the negotiation item that are taken into account during the evaluation. The idea is that it may be beneficial for people to introduce multiple issues in a negotiation when they have different preferences over these issues because it may be possible to trade off one issue for another in order to reach agreements where both the negotiators are better off. So, in multi-attribute negotiations the parties involved need to settle more than one issue. For example, agents may need to come to agreements that are characterized by attributes such as price, quality, delivery time, and so on. If the impact of the issues under negotiation over the satisfaction function is different for each agent (that is, some issues are more important for a participant than for the others and vice versa), the issues may be traded-off against one another, increasing the social welfare of the deal.

Both in single and multi-issue negotiations the outcome depends on four key factors (Fatima et al., 2006): the *negotiation protocol*, the *participant's strategies*, the *players' preferences* over the possible outcomes, and the *information that the participants have about the others*. However, in multi-issue negotiations appears an additional factor: the *negotiation procedure*, which specifies how the issues will be settled. There are three ways of negotiating multiple issues: *Package deal* which links all the issues and discusses them together as bundle, *Simultaneous negotiation* which settles the issues simultaneously, but independently, and *Sequential negotiation* which negotiates the issues sequentially one after another. This chapter will focus on a package deal based procedure.

As we pointed out before, many automated negotiation models have been developed. They may be classified regarding many different criteria (Buttner, 2006). Regarding their theoretical approach, *game theoretic*, *heuristic* and *argumentation-based* approaches exist. The game-theoretic approach tries to find optimal strategies by the analysis of the equilibrium conditions (Nash, 1950). Game-theoretic models are deemed mathematically elegant, but are very restricted in use because of their assumptions of unlimited resources, perfect rationality and a perfect information situation. In heuristic approaches the mentioned assumptions are relaxed, and players try to find an approximate solution strategy according to principles of bounded rationality by utilizing heuristic search and evaluation techniques (Faratin et al., 1998; Ehtamo et al., 1999; Faratin et al., 2002; Klein et al., 2003; Gatti & Amigoni, 2005; Lai et al., 2006; Ito et al., 2008;). Both in game-theoretic and heuristic approaches, negotiation protocols are usually based on the communication of offers in the form of potential agreements. In contrast, in argumentation-based negotiations, the agents are able to reason their positions including a meta-level component that may use promises, rewards, threats, as well as issue various forms of appeal (Rahwan et al., 2003). In addition to the theoretical approach criterion, negotiation can be classified regarding its structure, regarding the negotiation process, and regarding the restrictions over time and information situations.

In addition to the problem of selecting the optimal strategies in the negotiation processes, the agent's decision making mechanisms in multi-attribute negotiations have to face the problem of characterize the preference on all attributes. The characterization of preferences has a critical influence on the negotiation protocols and decision-making mechanisms. To end up with this introduction we briefly review some of the most relevant approaches in negotiation to model preferences, and pick up one of them to propose a multi-attribute negotiation protocol that will be presented in the following sections.

A typical way to model preferences is to use *utility functions*. In the case of multiple attributes, we talk about *multi-attribute utility theory (MAUT)*. Another approach to model preferences is to employ *multi-criteria decision making (MCDM)* (also called *multi-objective* or *multi-criteria optimization*) theory. In MCDM an agent has several objectives that are statements that delineate the desires of a decision maker. Thus, an agent wishes to maximise his objectives, which in some cases will conflict which each other in that the improved achievement with one objective can only be accomplished at the expense of another. Given an assignment of values to the corresponding attributes an agent measures how much the different objectives are fulfilled. Finally, a utility function is applied over the set of different levels of satisfaction of the agent's objectives. Research on those topics is conducted mostly in the field of decision theory. In the negotiation models described in the literature which use the utility based approaches to the modelling of preferences, the negotiation protocols are based on the communication of offers and counteroffers expressed as an assignment of values to the corresponding attributes. This approach to negotiation is known as *positional bargaining*, and is the predominant form of negotiation in the game-theoretic and heuristic approaches to negotiation. On the other hand, in argumentation-based negotiation the exchange of offers and counteroffers includes meta-information with the aim of reason the agents' positions. In the area of interest-based negotiation, another way to modelling preferences is to use *constraints* to restrict the attribute values that are preferred. *Constraints* in different formats, from *fuzzy* to *probabilistic* or *weighted constraints*, have been used in several models and approaches to multi-attribute negotiation (Luo et al., 2003; Lai & Lin, 2003; Ito et al., 2008). There are three main reasons that make very convenient the use of constraints as the core of a negotiation model. First, it is an efficient way of capturing requirements; second, constraints are capable of representing trade-offs between the different possible values for attributes; and third, using constraints to express offers in turns means that the solution space can be explored in a given exchange and so means that the search for an agreement is more efficient than in positional bargaining. The negotiation framework presented in this chapter falls within the heuristic approaches to non-mediated multi-attribute bilateral negotiations under incomplete information settings, and uses fuzzy constraints to model agent's preferences. With incomplete information we mean that agents lack information about other's discounting factors, reservation prices, utility functions or deadlines, and with non-mediated we mean that agents negotiate without the intervention of a mediating agent. The negotiation model is based on the hypothesis that by means of an expressive approach to constraint based negotiation the negotiation processes may be more efficient than with other approaches where mainly positional bargaining is used. Behind this is the idea that with the cost of a bounded increase in the revelation of private information, the decision mechanisms are more accurate when searching the negotiation space.

The remainder of this chapter is organized as follows. The next Section recalls the most relevant concepts on modelling agent's preferences and presents some preliminaries. Section 3 presents an example negotiation scenario where two different negotiation techniques are applied in order to show the possible advantages of expressive negotiation. Then the negotiation framework followed by an empirical evaluation is described. Finally, Section 6 presents the conclusions.

2. Modelling agent's preferences

A multi-attribute negotiation can be seen as a distributed *multi-objective optimization problem*. The participants in a negotiation have their own preferences over the negotiated attributes, and these preferences can be formulated in its most extensive form as a multi-objective or multi-criteria decision making problem. By definition, objectives are statements that delineate the desires of a decision maker. Thus, an agent wishes to maximise his objectives. However, it is quite likely that a decision maker's objectives will conflict with each other in that the improved achievement with one objective can only be accomplished at the expense of another. Therefore, a negotiator agent has to settle at a compromise solution. This is the topic of the *multi-criteria decision making theory*. Part of the solution to this problem is that the agent has to identify or approximate the *Pareto frontier* in the *consequence space* (i.e. in the space of the satisfaction levels of the different objectives). This task can be accomplished using different methods based on standard optimization techniques. Regarding the negotiation process it can be seen as a special case of multi-objective optimization problem. In this case, we have a set of distributed agent's objectives that should be satisfied. Each agent's objective depends on his individual objectives. The question now is if we can compute the Pareto frontier in a similar way. Assuming a set of agents which formalize their preferences as a multi-objective decision making problem, and that each agent computes his Pareto frontier, the only way to solve this problem in a similar way would be to share this information to formulate the global multi-objective optimization problem. In practice, this could be done by means of a trusted mediator, but it has a fundamental problem, agents and humans try to minimise the revelation of private information in negotiation to avoid strategic manipulation. Moreover, though Pareto optimality is a key concept in multi-objective optimization, we cannot forget that the aim of the negotiation is to reach an agreement, and so, it is necessary to pick up a fair solution from the Pareto frontier. However, fairness is not an easy concept to manage in negotiations.

2.1 Multi-attribute decision problems

As we stated before, negotiator agents are decision makers, and their decisions are based on preferences over the values of the different attributes. Formally, a *Multi-Attribute Decision Problem (MADP)* is defined as a set of attributes $X = \{x_1, ..., x_n\}$; a set of domain values $D = \{D_1, ..., D_n\}$ where each D_i is a set of possible values for attribute x_i; a set of constraints $C = \{C_1, ..., C_m\}$ where each C_j is a constraint function on a subset of attributes to restrict the values they can take; a set of available outcomes $O = \{o_1, ..., o_l\}$ where each o_j is an element of the possible outcome space *D*, and *O* is a subset of *D*; and a set of decision maker's preference statements $P = \{P_1, ..., P_m\}$. Agents negotiate over the same set of attributes and domain values, but each agent has a different set of constraints, available outcomes and preference statements. In a negotiation process, agents try to maximize their preferences, and in order to compute those values they have to solve the MADP. Among the different approaches to model agents' preferences from the MADPs perspective we survey two different categories of methods: the *constraint satisfaction problem (CSP) framework*, and the *multi-attribute utility theory (MAUT)*. For a detailed survey including more methods on MADPs see (Zhang & Pu, 2005).

A **CSP** is defined by a 3-tuple <*X,D,C*>, where *X* is a set of variables, *D* is a set of domains and *C* is a set of constraints. A solution to a CSP is a set of value assignment

$v = \{x_1 = v_1, ..., x_n = v_n\}$ where all constraints in *C* are satisfied. Therefore, the constraints are crisp (hard) since they are either respected or violated. A number of different approaches have been developed for solving this problem. One simple approach is to simply *generate-and-test*. However, when the CSP is complex the algorithm is not practical due to the computational complexity. A more efficient method is the *backtracking algorithm* that essentially performs a depth-first search of the space of potential CSP solutions. However, the complexity of backtracking for most nontrivial problems is still exponential. Other search algorithms for classical CSPs include: *forward checking*, *partial lookahead*, *full lookahead*, and *really full lookahead*.

We can see how a solution of a classical CSP needs to satisfy all the crisp constraints. Comparing the definition of classical CSP and MADP we can see that the main difference between them is that the MADP has a set of preferences, some of which can be violated when finding the optimal solution. Classical CSPs have been extended to *soft CSPs* in which not all the given constraints need to be satisfied. In the following, we recall several kinds of soft CSPs and a general framework which describes both classical and soft CSPs.

Fuzzy CSPs (FCSPs) extend the hard constraints by *fuzzy constraints*. A fuzzy constraint is a mapping from the direct product of the finite domain of the variables referred by the constraint to the [0,1] interval. The solution of a fuzzy CSP is the set of n-tuples of values that have the maximal value. The value associated with each n-tuple is obtained by minimizing the values of all its sub-tuples. An FCSP can be solved in a similar way as classical-CSP turning all fuzzy constraints into hard constraints.

Probabilistic CSPs (PCSPs) model those situations where each constraint *c* has a certain independent probability *p(c)* to be part of the given real problem. Let v be an n-tuple value set, considering all the constraints that the n-tuple violates, we can see that the probability of n-tuple being a solution is $\prod_{\text{all c that v violates}} (1 - p(c))$. The aim of solving PCSPs is to get the n-tuple with the maximal probability. The main difference between FCSPs lies in the fact that PCSPs contain crisp constraints with probability levels, while FCSPs contain non-crisp constraints. Moreover, the criteria for choosing the optimal solutions are different.

Weighted CSPs (WCSPs) allow to model optimization problems where the goal is to minimize the total cost of a solution. There is cost function for each constraint, and the total cost is defined by summing up the costs of each constraint. Usually WCSPs can be solved by the *Branch and Bound algorithm*.

A semiring-based CSP framework describes both classical and soft CSPs. In this framework, a semiring is a tuple *(A,+,x,0,1)* such that: *A* is a set and $0,1 \in A$; + is a close, commutative, and associative operation on A and 0 is its unit element; x is a closed, associative, multiplicative operation on A; and 1 is its unit element and 0 is its absorbing element. Moreover, x distributes over +. A c-semiring is a semiring such that + is idempotent, x is commutative, and 1 is the absorbing element of +.

Both the classical CSPs and the different type of soft CSPs can be seen as instances of the semiring CSP framework. The classical CSPs are Semiring-CSPs over the semiring $S_{CSP} = (\{false, true\}, \vee, \wedge, false, true)$ which means that there are just two preferences (*false* or *true*), that the preference of a solution is the logic *and* of the preferences of their subtuples in the constraints, and that *true* is better than *false*. FCSPs can be represented by $S_{FCSP} = ([0,1], \max, \min, 0, 1)$ which means that the preferences are over [0,1], and that we want to maximize the minimum preference over all the constraints. Similarly, the semiring corresponding to a PCSP is $S_{PCSP} = ([0,1], \max, \times, 0, 1)$, and the WCSPs can be represented by the semiring $S_{WCSP} = (R^+, \min, +, +\infty, 0)$.

Utility theory and MAUT has been used in solving decision problems in economics especially for those involving uncertainty and risk. Given the utility function, the decision maker's preferences will be totally determined, and the optimal solution will be the outcome with the maximal utility. When using MAUT to solve a multi-attribute decision problem that only involves certainty, the main task is to assess the value function according to the decision maker's preferences.

Let $O = \{O_1, ..., O_n\}$ be a set of outcomes of the MADP, $\mathfrak{A}$ be the set of all lotteries on the set O where $\sum p_i o_i \in \mathfrak{A}$, $p_i \in [0,1]$, and $\sum p_i = 1$; and $\succsim$ be a binary relation on $\mathfrak{A}$. First we define 4 axioms: 1) $\succsim$ is complete, i.e. either x $\succsim$ y or y $\succsim$ x; 2) $\succsim$ is transitive, i.e. if x $\succsim$ y and y $\succsim$ z, then x $\succsim$ z; 3) Continuity: given x $\succ$ y $\succ$ z, then there is an $\alpha, \beta \in (0,1)$ such that $\alpha x + (1-\alpha)z \succ y$ and $y \succ \beta x + (1-\beta)z$; 4) Independence: for all $x, y, z \in \mathfrak{A}$ and any $\alpha \in [0,1]$, x $\succsim$ y if and only if $\alpha x + (1-\alpha)z \succsim \alpha y + (1-\alpha)z$. Then the *von Neumann Morgenstern Theorem* proved the existence of utility function theoretically provided that the relation $\succsim$ satisfies the four axioms: Let $\mathfrak{A}$ be a convex subset of a linear space, and let $\succsim$ be a binary relation on $\mathfrak{A}$, then $\succsim$ satisfies the four axioms if and only if there is a real-valued function $u: \mathfrak{A} \rightarrow \mathfrak{R}$ such that:

a. $\forall x, y \in \mathfrak{A}, x \succsim y \Leftrightarrow u(x) \geq u(y)$;

b. $\forall x, y \in \mathfrak{A}$ and $\forall \alpha \in (0,1), u(\alpha x + (1-\alpha)y) = \alpha u(x) + (1-\alpha)u(y)$.

The function *u* is called the *utility function*.

Keeney and Raiffa (Keeney & Raiffa, 1976) extended the utility theory to the case of multi-attributes. Multi-attribute utility theory is concerned with the valuation of the consequences or outcomes of a decision maker's actions. For a decision problem where each action has a deterministic outcome, the decision maker needs only to express preferences among outcomes. The preference relation can be captured by an order-preserving, real-valued value function. Then, the optimal problem of the multi-attribute decision problem can be converted into the format of the standard optimization problem to maximize *u(x)*. When there is uncertainty involved in the decision problem, the outcomes are characterized by probabilities. It must be noted that a utility function is a value function, but a value function is not necessarily a utility function. In the case that only certainty is involved, the utility and value function are interchangeable.

3. A non-mediated bilateral negotiation model based on fuzzy constraints

Here we propose a non-mediated fuzzy constraint based negotiation framework for competitive e-marketplaces in which multiple buyer agents negotiate bilaterally with multiple seller agents to acquire products. In competitive markets, there is an inherent need to restrict the amount of private information the agent reveals. However, this restriction can have a detrimental effect on the search for a solution. As we stated above, especially in the case of multi-attribute negotiations, it is possible to reach a more satisfactory agreement by means of an adequate combination of attributes or constraints. However, most solutions put forward to tackle this problem are mediated, iterative and approach mechanisms, which are applicable to preference models based on linear-additive or quasi-concave utility functions (Ehtamo et al., 1999; Faratin et al., 2002; Lai et al., 2006). Other approaches based on non-linear utility spaces include a mediator in the negotiation processes (Klein et al., 2003; Gatti & Amigoni, 2005; Ito et al., 2008). As an alternative to these solutions, we propose one based on the concept of communicative rationality rather than one which is merely strategic and retains as the fundamental criteria the minimization of private information revealed. Our solution is therefore based on a dialogue of offers in which preferences or satisfaction degrees are partially disclosed. The hypotheses on which the work is based is that of an interactive model which is sufficiently expressive to allow a discussion of proposals by means of a partial declaration of preferences which permits the agents to reach a more satisfactory agreement, being confined to the need to minimize the loss of privacy. The negotiation framework is defined by: a fuzzy constraint based model of preferences; the expressive behaviors and strategies of the agents; an interaction model that permits the automatic generation of expressive or non-expressive dialogues with different degrees of symmetry; and finally a set of decision mechanisms adapted to the interaction model and the preferences of the agents.

There are several works using fuzzy constraints to model preferences, however, most of them use single point offers (i.e. positional bargaining). The FeNAs (Fuzzy e-Negotiation Agent system) platform (Kowalczyk & Bui, 2000) uses fuzzy constraints and permits correlated multiple bilateral negotiations. It is one of the first works in which the problem of multi-attribute negotiation is clearly presented using a preference model based on FCSP. The main problem with FeNAs resides in its being a positional approach. Lai (Lai & Lin, 2004) presents a general framework for multi-attribute and multilateral negotiation based on fuzzy constraints. The negotiation model is based on FCSP, which when applied to a distributed domain of agents is organized as a network of distributed fuzzy constraints (DFCN). This work makes some very important contributions to the regularization of the mechanisms for calculating the satisfaction degree and to the available concession and compensation strategies. It introduces fuzzy logic techniques to the relaxation decision making area that allow concession strategies to be defined that are a function of the beliefs and desires of the agents. The model is also based on single-point offers and there is no argumentation, but decision-making is based on the behavior of the opponent and the type of offers received. In accordance with the mentioned above procedures, if there is no convergence in the first relaxation steps, the number of offers increases exponentially. If there are a large number of attributes, the number of possible proposals for a particular cut level becomes intractable. Although the similarity function can help with convergence, a certain amount of knowledge of the utility functions of the opponent is assumed. Finally, Luo (Luo et al., 2003) develops a fuzzy constraint based model for bilateral multi-issue negotiations in semi-competitive environments. It uses crisp constraints to express offers and includes the idea of rewards and restrictions. The most noticeable aspects are related to the acceptability function and with the operators used to apply the prioritization of the fuzzy constraints. Assuming the seller agents' dominant strategy is to offer the first product that satisfies the constraints, the model isn't efficient enough because it exhibits a large lack of symmetry. In this model a buyer agent has a great communication power (expressing offers by means of constraints) while the seller agent can only offer specific products or request a relaxation of the constraints. In this way, the opportunity to apply some form of solution compensation technique so that a win-win solution is obtained is lost.

3.1 Expressive vs inexpressive negotiation dialogues

In this subsection a bilateral negotiation scenario is presented, comparing two approaches, one expressive and the other non-expressive, in which all the advantages that our approach contributes to the problem will be discussed.

A buyer agent and a seller agent begin a negotiation dialogue about the sale of a vehicle. The buyer agent expresses a desire to buy in the following way: "*I want to acquire a car at a low price, of high quality and as new as possible*". From this statement, it can be taken there are three issues that are of interest to the buyer agent, the *price*, the *quality* and the *age* of the car. The requirements of the buyer agent are therefore defined by these three fuzzy constraints, so that a priori, no specific range is defined for each issue to determine whether a constraint has been satisfied. In the seller agent's case we could propose a formulation of preferences or sale needs in a similar way, however, in trading scenarios the seller agent may be more inclined towards the use of catalogues of products. In Figures 1 and 2, the buyer agent's preferences and a summary of the seller agent's catalogue are shown respectively. The labels above each step represent the range of the attributes value domain, in such a way that the states can appear as intervals, numeric groups or as linguistic terms. The higher steps represent greater satisfaction degrees. If we analyse the diagram we can see that, for

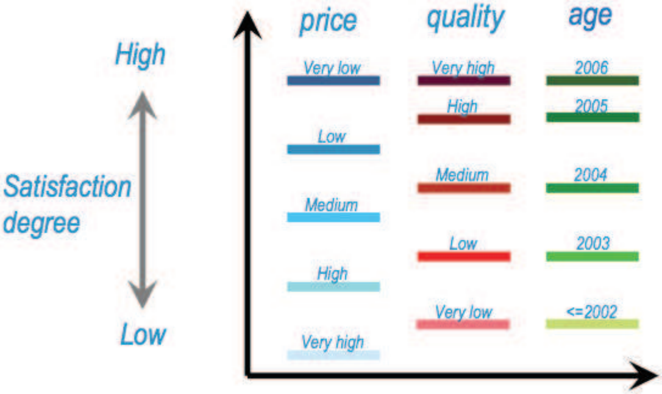

Fig. 1. Buyer agent's preferences

Fig. 2. Seller's catalogue of products

example, the fuzzy constraint expressed as low price is divided in intervals in accordance with the different satisfaction degrees of the buyer agent. The catalogue of products is defined by a series of rows each one of which characterizes a product. For each product, the satisfied range of values of the buyer agent's attributes is shown. The last column represents the utility the seller agent obtains if the product is sold. This utility value does not have to have any direct correlation with the negotiable attributes, there may exist other private issues (non negotiated) that have a greater influence on the utility value.

To give an example of our working hypothesis we first present a possible negotiation dialogue between a buyer and seller agent (see Figure 3) that we will call *non-expressive*. In this type of dialogue the argumentation capability with respect to the offers is minimal. The buyer agent makes offers in the form of crisp constraints taken strategically from the fuzzy constraints that represent its overall requirements. On the other hand, the seller agent is only able to accept or reject an offer. So, we see in the example that the buyer agent successively relaxes its demands, as after each offer the seller agent responds with a refusal (as it does not have products that satisfy the constraints). Finally, in the last stage, the seller agent finds a product p4 that satisfies the buyer agent's requirements. However, this solution provides a very low profit for the seller agent. It is clear that the negotiating position of the buyer agent is much stronger, their requirements are described in detail in each offer, and at no time does the seller agent give any clue as to its preferences. The limitations of the language used mean that the only possible criteria that can be used to find solutions are local preferences. The question we must ask ourselves is whether there exists a solution that would have been more satisfying for the seller agent without worsening the

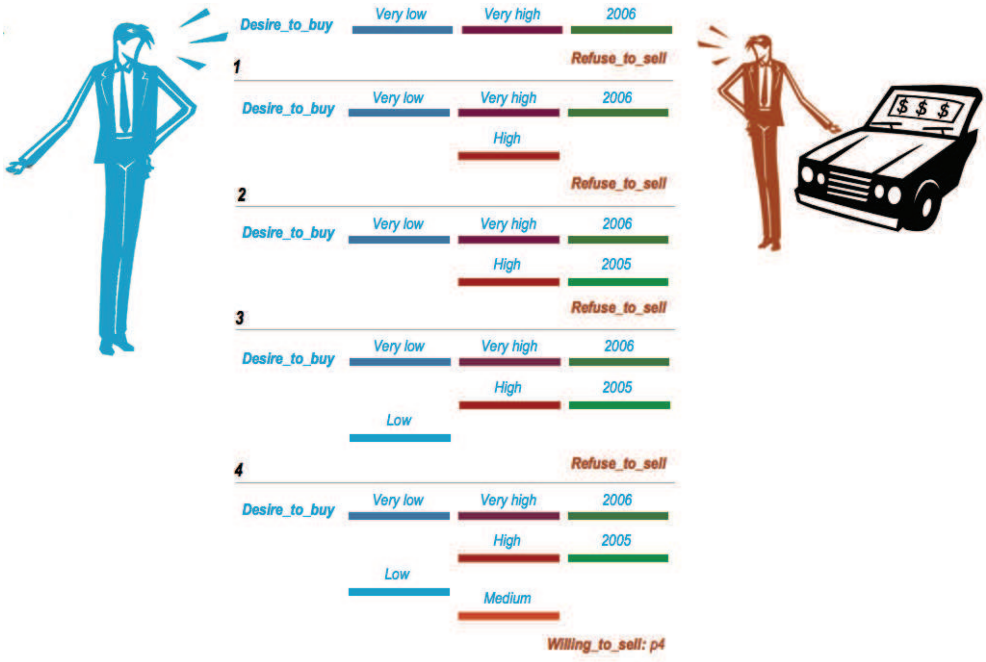

Fig. 3. Example of non-expressive dialogue

buyer agent satisfaction degree, and the answer rests in the solution p3, which would indeed have been more satisfactory for the seller agent without being less so for the buyer agent. As an alternative, we now present a new dialogue, which we term *expressive*, in which the concepts that form the basis of our hypothesis are applied. In Figure 4, the buyer agent and the seller agent negotiate the purchase of an automobile under the same preference conditions used in the previous dialogue. In this dialogue two important innovations appear: Firstly, the buyer agent is able to subjectively value its offers; and secondly, the seller agent is able to clarify its refusal to offer a product, by using expressions that allow it to state which constraints it wants the buyer agent to relax.

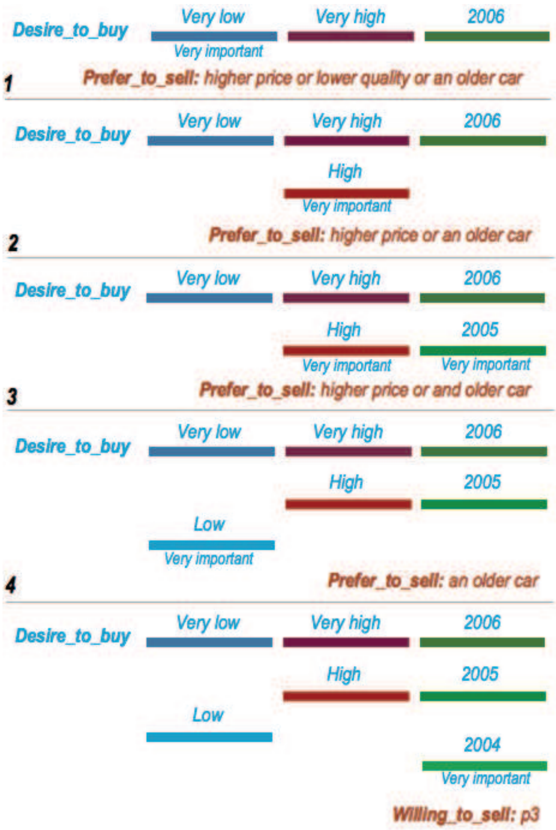

Fig. 4. Example of expressive dialogue

We will now analyse the course of the dialogue.

1. The first offer made by the buyer agent is the one that subjectively offers it the greatest satisfaction. Apart from the offer, defined as a set of crisp constraints, these constraints contain meta-information that grades them depending on the degree of importance each of them has. Thus, the constraint Very Low is considered as very important and it is expressed like this in the dialogue. The seller agent does not have a product that satisfies all the constraints, so it has no choice but to refuse the offer. However, it argues

its refusal with an attack based on preferences, suggesting that the buyer agent relax constraints with differing degrees of preference. From the seller agent point of view, any of the constraints in the initial dialogue can be relaxed.

2. The buyer agent's second offer involves relaxing the quality constraint. As the seller agent had no preference for which constraint should be relaxed, the buyer agent relaxes at random one of the constraints (quality or age) that least affects its satisfaction degree. The quality constraint now becomes the buyer agent's choice, because to do so later would involve a greater loss of satisfaction than the relaxation of any other constraint. When the seller agent receives the offer, it is unable to find a product that satisfies all the constraints. However, it concludes that products p2 and p3 come close to the buyer agent's requirements. To be precise, the seller agent reasons in the following way: p2 will provide me with more profit, but on the other hand, although p3 will provide me with slightly less profit, it is closer to the buyer agent's requirements. After the seller agent has made the previous reasoning, it tries to persuade the seller agent by first asking it to relax the price and age constraints.
3. The third stage of the negotiation follows similar parameters to the previous one.
4. In the buyer agent's fourth offer, the price constraint is the most important. The seller agent analyses its catalogue and rejects p1 because of its low utility and estimated distance. With regards to p2, it decides that it satisfies the age and quality constraints, and that p3 satisfies the price and quality constraints, and finally, that p4 satisfies the price and age constraints. A priori, the three products are relatively close to the buyer agent requirements, but the description of the price constraint as very important affects the estimation of the closeness or distance of p2. The distance of products p3 and p4 is estimated to be similar, so the buyer agent discriminates depending on the utility of the solutions. The conclusion is that the seller agent decides that p3 is the best possible offer. He then puts all its effort into ensuring the sale of p3, although it does not satisfy the age constraint, which is why the request to relax concentrates on this constraint.
5. After receiving the request to relax, the buyer agent finds that a priori, it has no problem with relaxing either the quality or the age constraint. Under the assumption of negotiation based on interests or principles, the buyer agent accepts the request to relax the age constraint. The seller agent has a product, p3 that satisfies the present requirements. The overall satisfaction of the solution is greater than in the case with non-expressive negotiation.

The challenge of developing all the concepts in the example involves several aspects. Firstly, an agents' preference model formalization. Secondly, a definition of the negotiation profile for modelling the agent's behaviour towards their opponents. Creation of a communication model that, amongst other things, details the locutions needed to be able to deal with all the expressive nuances. Development of decision making mechanisms. Finally, a working language specification allowing the decision mechanisms to be linked to the expressions available to the agents.

4. Negotiation framework

The negotiation framework consists of a description of the *agent's domain knowledge*; a *dialogue model*; the *decision mechanisms*; and the *transition rules* that connect the locutions to the mechanisms.

4.1 Agent's domain knowledge

Buyer agent's requirements over the attributes of a product are described by means of a *fuzzy constraint satisfaction problem (FCSP)*, which is a 3-tuple (X, D, C^f) where $X = \{x_i | = 1,...,n\}$ is a finite set of variables, $D = \{d_i | = 1,...,n\}$ is the set of finite domains of the variables, and $C^f = \{R_j^f \mid j = 1,...,m\}$ is a set of fuzzy constraints over the variables. It is worth noting that a fuzzy constraint may restrict more that one variable or attribute. A fuzzy constraint corresponds to the membership function of a fuzzy set. The function that numerically indicates how well a given constraint is satisfied is the satisfaction degree function $\mu_{R_j^f} : X \to [0,1]$, where 1 indicates completely satisfied and 0 indicates not satisfied at all. Given the *cut level* $\sigma \in [0,1]$, the induced crisp constraint of a fuzzy constraint R^f is defined as R^c. It simply means that if R^c is satisfied, the satisfaction degree for the corresponding fuzzy constraint will be at least σ. Therefore, the *overall (global) satisfaction degree (osd)* of a given solution $x' = (x_1', ..., x_n')$ is:

$$\alpha(x') = \min\{\mu_{R^f}(x') \mid R^f \in C^f\} \tag{1}$$

On the other hand, a seller agent owns a private catalogue of products $S = \{s_k \mid s_k = (p_k, u_k)\}$, where p_k is the vector of attributes and u_k is the profit the seller agent obtains if the product is sold. We assume that the profit u_k may depend not only on the negotiated attributes but also on non-negotiated ones (stock period for instance).

Let A_b and A_s represent a buyer and a seller agent, a negotiation process is a finite sequence of alternate proposals from one agent to the other. During the negotiation stage, A_b utters *purchase requirements*,

$$\pi = \bigcap \left\{ R_j^{c(\sigma_j)} \mid j \in [1, m] \right\} \tag{2}$$

where $R_j^{c(\sigma_j)}$ is a crisp constraint induced from R_j^f at a cut level σ. Therefore, a purchase requirement is a purchase proposal that is formed by a set of crisp constraints extracted from the set of fuzzy constraints that describes the buyer's preferences regarding the attributes of the products. Each crisp constraint in the purchase requirement can be induced at a different cut level. Complementing the *osd* definition, the *potential or expected overall satisfaction degree (posd)* is the osd that a buyer agent may get if the corresponding purchase requirement is satisfied. It is defined as:

$$\alpha^{\pi} = \min\{\sigma_i \mid i = 1, ..., m\} \tag{3}$$

A seller agent may respond to a buyer agent in three different ways: *rejecting the proposal*, *offering a product* that satisfies the purchase requirement, or *suggesting the relaxation* of the purchase requirement. A *relaxation requirement* is defined as a set:

$$\rho = \{r_j \mid r_j \in [0,1]\} \tag{4}$$

where ρ_j is the preference for constraint *j* to be relaxed. The negotiation process and the agreements achieved will mainly vary depending on the strategies followed by the agents when generating purchase requirements and when requesting its relaxation. We cover all these aspects modeling the *agents' attitudes*. Agents' attitudes are related to the agents' strategic behavior in the negotiation process, where strategic behaviors are described in terms or expressiveness and receptiveness. A negotiation profile $\Pr ofile_{\text{seller}} = \{\psi, \beta\}$ describes the seller agent's attitude, where $\psi \in \{0,1\}$ controls whether it uses or not relaxation requests in order to express its preferences for a specific relaxation of the previous buyer's demands, and $\beta \in [0,1]$ modulates its attitude regarding a purchase requirement received from a buyer agent. Finally, a negotiation profile $\Pr ofile_{\text{buyer}} = \{\xi, \eta\}$ describes the buyer agent's attitude, where $\xi \in \{0,1\}$ controls whether it uses or not *purchase requirement valuations* defined as:

$$v = \{v_j \mid v_j \in [0,1]\} \tag{5}$$

where v_j is the degree of importance that the constraint j has for the buyer agent, and $\eta \in [0,1]$ modulates its attitude regarding a relaxation requirement received from a seller agent.

4.2 Negotiation dialogue

The framework of *formal dialogue games* is increasingly used as a base for structuring the interactions of agents communication protocols (McBurney et al., 2003), adopted from the theory of argumentation field. Formal dialogue games are those in which two or more players pronounce or transmit locutions in accordance with certain predetermined rules. In our negotiation model all dialogues are confined to two agents, one the buyer and the other the seller, so that the dialogues are exclusively bilateral. A dialogue is structured in accordance with the following stages:

1. *Opening* the dialogue.
2. *Negotiation*: this stage is defined by a sequence of iterations that are based on the domain knowledge mentioned earlier. These iterations are now itemised:
   - Buyer agent:

- Transmit purchase requirements.
- Transmit valuation of purchase requirements.
- Reject sale offers.
   - Seller agent:
- Transmit sale offers.
- Rejects purchase requirements.
- Propose the relaxation of purchase requirements.
- Reject purchase obligations.

3. *Confirmation*: the participants come to a compromise and reach an agreement.
4. *Close* of dialogue: the dialogue ends.

Our dialogue proposal is subject to the following rules:

a. The first stage in the dialogue is Opening of the dialogue.
b. The Opening and Closing stages of the dialogue can only occur once in the whole dialogue.
c. The only stages that must appear in all dialogues that end normally are Opening and Closing of the dialogue.
d. The Confirmation stage requires the negotiation stage to have occurred previously.
e. The last stage of all dialogues that end normally is Close of dialogue.

The participants can commute between the negotiation and confirmation stages, subject only to the rules and the constraints defined by the combination of locutions rules, which we describe later.
Our purchase negotiation dialogue is defined as sequence of four stages: **open dialogue** (L1-2), **negotiate** (L3-8), **confirm** (L9-10) and **close dialogue** (L11).

L1: open_dialogue (P_b, P_s, θ) P_b suggests the opening of a purchase dialogue to a seller participant P_s on product category θ. P_s wishing to participate must respond with *enter_dialogue*(.).

L2: enter_dialogue (P_s, P_b, θ) P_s indicates a willingness to join a purchase dialogue with participant P_b. Within the dialogue, a participant P_b must have uttered the locution *open_dialogue*(.).

L3: willing_to_sell (P_s, P_b, p_j) P_s indicates to the buyer P_b a willingness to sell a product. A buyer P_b must have uttered a *desire_to_buy*(.) or a *prefer_to_buy*(.) locution.

L4: desire_to_buy $(P_b, P_s, \pi_{\mathcal{B}_{req}})$ P_b, speaking to the seller P_s, requests to purchase a product that satisfies the purchase requirement $\pi_{\mathcal{B}_{req}}$.

L5: prefer_to_sell $(P_s, P_b, \pi_{\mathcal{B}_{req}}, \rho_{\mathcal{B}_{req}})$ P_s, speaking to the buyer, requests to relax the purchase requirement $\pi_{\mathcal{B}_{req}}$, and expresses which constraints are preferred to be relaxed, by means of the relax requirement $\rho_{\mathcal{B}_{req}}$.

L6: prefer_to_buy $(P_b, P_s, \pi^k_{\mathcal{B}_{req}}, \upsilon_{\mathcal{B}_{req}})$ P_b, speaking to the seller, requests to purchase a product which satisfies the purchase requirement $\pi^k_{\mathcal{B}_{req}}$, and expresses its preferences for the different constraints by means of the purchase requirement valuation $\upsilon_{\mathcal{B}_{req}}$.

L7: refuse_to_buy (P_b, P_s, p_j) Buyer agent expresses a refusal to purchase a product. This locution cannot be uttered following a valid utterance of *agree_to_buy*(.).

L8: refuse_to_sell $(P_s, P_b, p_j \mid \pi_{\mathcal{B}_{req}})$ Seller agent expresses a refusal to sell a product, or it expresses a refusal to sell products that satisfy the purchase requirement $\pi_{\mathcal{B}_{req}}$. This locution cannot be uttered following a valid utterance of *agree_to_sell*(.).

L9: agree_to_buy (P_b, P_s, p_j) Buyer agent P_b speaking to P_s commits to buy a product. A locution of the form *willing_to_sell*(.) must have been uttered.

L10: agree_to_sell (P_s, P_b, p_j) Seller agent speaking to buyer agent commits to sell a product. A locution of the form *agree_to_buy*(.) must have been uttered.

L11: withdraw_dialogue (P_x, P_y, θ) For P_x and P_y participants with different roles (i.e. sellers and buyers). P_x announces agent P_y the withdrawal from the dialogue.

Next step is to specify the mechanisms that will invoke particular locutions in the course of a dialogue.

4.3 Decision mechanisms

Syntactic rules are not enough to ensure that the dialogues are generated automatically. It is essential to equip each participant with mechanisms that allow it to invoke the correct locution at the right time, as a response to previous locutions or in anticipation of future ones. This type of mechanism we term *semantic decision mechanism*. The mechanisms are grouped together depending on the role of the participant: Buyer (B) or Seller (S). We will now describe each mechanism's general directive and then detail their specific features. In addition, we specify the output generated by the mechanisms, a key point for describing, in the following Section, the working features or working semantics that connect the decision mechanisms and the locutions. We begin with the buyer agent's decision mechanisms.

B1: Recognize Need Allows a buyer agent to recognize the need to acquire a product. This recognition may be as a consequence of the explicit initiative of the user (e.g. through an interface the user gives an order to their personal agent of their intention to acquire a product), or it could be an automatic response based on thresholds that are triggered automatically (e.g. when a personal agent detects that it is within range of an electronic market that offers a particular type of product that falls within the preferences of the owner of the personal agent). When it detects the need and furthermore interprets that it is possible to begin a dialogue the mechanism's output is *have_need* (θ). Outputs: *wait*, *have_need* (θ), *have_no_need* (θ), where (θ) defines a *product category*.

B2: Generate Purchase Requirement This mechanism responds to a buyer agent's need to generate purchase requirements. Any purchase requirement must be compatible with the locution *desire_to_buy*(.) or *prefer_to_buy*(.). Two possible outputs are recognized, one that states that it is impossible to generate a requirement and another that specifies the requirement generated. Outputs: *empty_set* $\varnothing$, $\pi_{\mathcal{B}_{req}}$

The method for extracting crisp constraints directly affects the way a purchase requirement is accepted, and indirectly affects the potential overall satisfaction degree the buyer agent hopes to obtain. There are two possible strategies when extracting crisp constraints to satisfy a purchase requirement and generate a specific potential overall satisfaction degree:

(Concession strategy) Given a purchase requirement $\pi^t_{\mathcal{B}_{req}}$ sent at an instant $t \in \mathbb{N}$, a general concession strategy is defined as mechanism that generates a new purchase requirement $\pi^{t+1}_{\mathcal{B}_{req}}$ so that $\alpha^{\pi^{t+1}_{\mathcal{B}_{req}}} < \alpha^{\pi^{t}_{\mathcal{B}_{req}}}$ and $\alpha^{\pi^{t+1}_{\mathcal{B}_{req}}} \geq \alpha^{\pi^{t}_{\mathcal{B}_{req}}} - \varepsilon$, where $\varepsilon \in (0,1]$.

According to this definition, ε is an arbitrary value that fixes the maximum loss of potential overall satisfaction that the buyer agent is willing to accept when generating a new purchase requirement. It determines the agent’s behaviour with respect to how rapidly it is willing to make concessions over its purchase requirements.

(Compensation strategy) Given a purchase requirement $\pi^t_{\mathcal{B}_{req}}$ sent in an instant $t \in \mathbb{N}$, a compensation strategy is a mechanism that generates a new purchase requirement $\pi^{t+1}_{\mathcal{B}_{req}}$ so that $\alpha^{\pi^{t+1}_{\mathcal{B}_{req}}} \geq \alpha^{\pi^{t}_{\mathcal{B}_{req}}}$.

We now move on to the specific mechanisms that put these strategies into practice. There are two ways to generate a new purchase requirement:

Adding a new fuzzy constraint. This way of generating purchase requirements is intended for two specific situations: the start of the negotiation, when the first purchase requirement should be prepared, and during the negotiation, after a sale offer that does not satisfy the constraints not included in the purchase requirement. In the first case, the agent selects a fuzzy constraint and applies the highest cut level to extract the corresponding crisp constraint and create the purchase requirement $\pi^t_{\mathcal{B}_{req}}$. By using this method, the agent is following the minimum revelation of information principle and the requirement obtained generates the greatest potential overall satisfaction degree. In the second case, a new constraint is selected from amongst those not satisfied by the sale offer received.

Modification of a previous purchase requirement. This way of creating a purchase requirement is intended for a specific situation: the locutions *prefer_to_sell*(.) or *refuse_to_sell*(.) sent by the seller agent during the negotiation with the intention of expressing its refusal to satisfy the buyer agent's requirements. Given a purchase requirement $\pi^t_{\mathcal{B}_{req}}$, and after receiving one of these locutions as a reply, the cut levels associated with the fuzzy constraints included must be changed and this change affects the potential overall satisfaction degree. Therefore, the generation of a new requirement $\pi^{t+1}_{\mathcal{B}_{req}}$ will be the aim of the application of one of the concession or compensation strategies and the problem is reduced to determining the plan for relaxing the previous purchase requirement $\pi^t_{\mathcal{B}_{req}}$. We propose the meta-strategy, which consists, when possible, of applying the compensation method and in its absence the concession method. The following algorithm implements the required function.

Algorithm 1. (*Modification of purchase requirements*)

1. Given a purchase requirement $\pi^t_{\mathcal{B}_{req}}$, a vector is obtained with the potential overall satisfaction degrees for all the possible purchase requirements resulting from the relaxation each time of only one of the constraints contained in $\pi^t_{\mathcal{B}_{req}}$:

$$[\alpha^{\pi^{(t+1)k_1}_{\mathcal{B}_{req}}} ... \alpha^{\pi^{(t+1)k_i}_{\mathcal{B}_{req}}}]$$

   where $\alpha^{\pi^{(t+1)k_x}_{\mathcal{B}_{req}}}$ represents the potential overall satisfaction degree obtained if the constraint $R^f_{k_x}$ is relaxed the minimum possible. The constraints that cannot be relaxed must be eliminated from the vector. If none of the constraints can be relaxed the function returns $\varnothing$.

2. The maximum of the previous vector is calculated:

$$\alpha^{t+1}_{max} = max([\alpha^{\pi^{(t+1)k_1}_{\mathcal{B}_{req}}} ... \alpha^{\pi^{(t+1)k_i}_{\mathcal{B}_{req}}}])$$

3. A new vector $\overrightarrow{\alpha^{t+1}_{max}}$ is generated in which only the elements that satisfy the following equality are included:

$$\alpha^{t+1}_{max} = \alpha^{\pi^{(t+1)k_x}_{\mathcal{B}_{req}}}$$

4. Finally, the following function is applied:

$$\arg\max_{\overrightarrow{\alpha^{t+1}_{max}}, \rho^t_{max}} \alpha^{\pi^{(t+1)k_x}_{\mathcal{B}_{req}}} + r_{k_x} * \eta$$

where ρ^t_{max} is a relax requirement from the seller agent, in which only those constraints included in the vector created in stage 3 are taken into account. If there are no relax requirements, r_{k_x} always takes the value 0. This function selects the constraint or constraints that maximize the total potential overall satisfaction that is induced if they are relaxed and of the relax requirement correspondingly weighted by the value η of the buyer agents receptive profile.

5. Once the constraint or constraints with the option of being relaxed are selected, one is chosen and a new purchase requirement is created $\pi_{\mathcal{B}_{req}}^{t+1}$, in which only the chosen constraint is modified.

The first three stages of the algorithm focus on the search for those constraints in $\pi_{\mathcal{B}_{req}}^{t}$ which if relaxed involve the smallest possible loss of potential overall satisfaction. Once these constraints have been detected, stage 4) takes into account only these constraints and, if there is a relax requirement, what the seller agent's preferences are in this respect. At one extreme if $\eta = 0$, the only criteria for relaxing is local, whereas if $\eta = 1$, the maximum importance possible is being given to the seller agent's recommendations. It is important to clarify that as we have defined in stage 2) the maximum value for filtering the potential satisfaction values, the function defined in 4) would only vary if r_{k_x} varies, being $\alpha^{\pi_{\mathcal{B}_{req}}^{(t+1)k_x}}$ a constant the same as α_{max}^{t+1}. However, we have decided to show the function in a more general form so we can easily extend the criteria of the maximum to other criteria.

B3: Generate Purchase Requirement Valuation This mechanism allows a valuation argument to be generated for a purchase requirement that has not yet been sent, i.e. a purchase requirement valuation $\upsilon_{\mathcal{B}_{req}}$. This can be communicated by the locution *prefer_to_buy*(.). The impossibility of obtaining a valuation generates the output an *empty_set*. Taking into account that the argumentation of a requirement is a reflection of the expressive character of the buyer agent, the mechanism will be controlled by its *expressive profile* ξ. If this has the value 1 the mechanism activates and tries to generate the valuation, if it has the value 0, the mechanism does not activate a valuation and returns an *empty_set*. When there are no valuations the buyer agent uses the locution *desire_to_buy*(.), whereas if there are valuations it uses the locution *prefer_to_buy*(.). Outputs: *empty_set* $\varnothing$, $\upsilon_{\mathcal{B}_{req}}$

A valuation of a purchase requirement is an expression of how important for the buyer agent the satisfaction of each of the purchase requirements constraints is. We propose the following algorithm.

Algorithm 2. (*Valuation of a purchase requirement*)

1. Given a purchase requirement, by sending $\pi_{\mathcal{B}_{req}}^{t+1}$, a vector is obtained with the potential overall satisfaction degrees for all the possible purchase requirements that result from relaxing only one of the constraints in $\pi_{\mathcal{B}_{req}}^{t+1}$ each time. The potential overall satisfaction degrees of those constraints that cannot be relaxed have the value 0:

$$[\alpha^{\pi_{\mathcal{B}_{req}}^{(t+2)k_1}} \ldots \alpha^{\pi_{\mathcal{B}_{req}}^{(t+2)k_i}}]$$

2. The elements of the previous vector are taken and a new standardized vector is defined that represents the valuation of the purchase requirement:

$$\upsilon_{\mathcal{B}_{req}} = [1-\alpha^{\pi_{\mathcal{B}_{req}}^{(t+2)k_1}} \ldots 1-\alpha^{\pi_{\mathcal{B}_{req}}^{(t+2)k_i}}] / sum([1-\alpha^{\pi_{\mathcal{B}_{req}}^{(t+2)k_1}} \ldots 1-\alpha^{\pi_{\mathcal{B}_{req}}^{(t+2)k_i}}])$$

The mechanism defines the valuation strategy of the purchase requirement as a strategy based on potential satisfaction degrees. To clarify which potential satisfaction degrees we are talking about we will describe a normal valuation process. When mechanism *B2:*