Coping with One-to-Many Multi-Criteria Negotiations in Electronic Markets

Bartel Van de Walle and Sven Heitsch
Computer and Information Science Department
New Jersey Institute of Technology
{bartel,sh2}@njit.edu

Peyman Faratin
Center for Coordination Science
MIT Sloan School of Management
peyman@mit.edu

Abstract

We are concerned with negotiators who negotiate on issues (products or services) characterized in terms of multiple criteria. In a typical round of the negotiation, the negotiator receives from various counter-parties responses or offers with specific values on the criteria that characterize the issue at stake. Clearly, the negotiator has certain preferences with respect to these offers. The analyis we present here allows us to reveal the dependencies among the various offers in light of the negotiator's preferences. In addition, we are able to expose the dependencies among the criteria, as they are implied by the counter-parties' offers. We make use of basic concepts from preference modeling and relational analysis to obtain a family of partial rank orders illustrating the dependencies among offers and criteria. We believe that, by using this information, the negotiator is better armed to identify the most suitable counter-party to negotiate with, or the criterion on which to focus in this negotiation.

1. Introduction

We are interested in developing a quantitative model of decision making for multi-lateral negotiation over multiple criteria. This class of problems is best seen in the Request For Quote (RFQ) protocol that is becoming increasingly popular for electronic procurement systems [2]. Agents can also implement such a quantitative model as decision algorithms to (semi)-autonomously negotiate the procurement of goods or services [7].

To focus our minds, assume that a negotiation takes place among one seller and multiple buyers. The seller has registered a particular product on an electronic marketplace, and provided a description in terms of multiple criteria that characterize that product, some of which are negotiable (for instance price and date of availability). She then has received various proposals or offers from interested buyers, each of which contains values on the indicated negotiable criteria. As every offer reflects a buyer's personal valuation of the product, it is not very likely that all offers will be identical. The obvious problem the seller then must solve is: given these offers, which buyer should I focus on to negotiate further and sell my product [6]?

One common solution is to calculate the total value of every potential buyer using a weighted sum multi-criteria model, and then rank order these values. The preferred buyer then is the buyer with the highest total value. While this method is straightforward, at least in principle, it severely limits the seller's understanding of the offers of the buyers. Indeed, the entire offer of every buyer is reduced to a single number – the total value.

We propose a second solution, taking into account all information contained within the buyers' offers, and preserving that information richness throughout the analysis. Based on pairwise comparisons of the seller's preferences on the buyers' offers, we again obtain a rank order of the different buyers. However, in contrast to the linear ordering offered by the first approach, the rank order is now only partial: two buyers can be incomparable, due to conflicting preferences on some criteria. Incomparability is a key relation in preference modeling [4], and adds substantial interpretational value to the standard analysis based on strict preferences and indifferences only.

We introduce the theoretical framework of our analysis, preference modeling and relational analysis, in the following section. In Section 3, we provide a detailed numerical illustration of our approach to the analysis of a seller negotiating with multiple buyers. We first illustrate how to obtain a family of partial rank orders for the buyers, and then we turn to the analysis of the criteria that characterize the issue at stake. Finally, we summarize our approach and point at our future research directions in Section 4.

2. Preference modeling

2.1. Crisp and fuzzy preference relations

A relation *R* from a set X to a set Y is a rule that assigns certain objects in the set X to certain other objects in the set Y. If both sets X and Y are identical, we say that *R* is a binary relation *on* X. The expression "*x* is in relation *R* with *y*" (or also "*x* is *R*- related to *y*") is denoted as *xRy* or $(x,y)\in R$. In preference modeling, the objects we are looking at are decision alternatives (or in short alternatives), and the decision rules express a decision maker's preference (or lack thereof, either an indifference or an incomparability) among all possible pairs of alternatives. Denoting the set of alternatives as *A*, then the relations expressing such preferences (or indifference or incomparabilities) all are binary relations in *A*. More specifically, we define the following three important relations on *A*:

- A couple of alternatives (*a*,*b*) belongs to the *strict preference* relation *P* if and only if the decision maker prefers *a* to *b*;
- A couple of alternatives (*a*,*b*) belongs to the *indifference* relation *I* if and only if the decision maker is indifferent between alternatives *a* and *b*;
- A couple of alternatives (*a*,*b*) belongs to the *incomparability* relation *J* if and only if the decision maker is unable to compare *a* and *b*, for instance caused by conflicting or insufficient information.

There are lots of properties for which it is difficult to obtain a crisp partition of the universe of objects into those objects that satisfy the property and those objects that do not satisfy the property. In preference modeling, for instance, it is usually not obvious to decide which decision alternatives are unequivocally preferred to what other alternatives. A solution to this crisp dichotomy is to introduce a general transition from membership (e.g., definitely preferred) to non-membership (e.g., definitely not preferred), and allowing for partial degrees of membership. Mathematically, this idea is translated into a fuzzy set.

Formally, a fuzzy set *A* on a universe X is a mapping from X to the unit interval [0,1], with the value *A*(x) of *A* in *x* of X the degree of membership of *x* in *A*. $A(x)=1$ means full membership, $A(x)=0$ means non-membership, and all values *A*(*x*) in]0,1[denote partial membership.

In order to analyze a fuzzy set *A* in X at a particular membership degree $\alpha \in [0,1]$, we can 'cut' the fuzzy set at that degree, and consider only the set of elements *x* of X that have a membership degree *A*(*x*) of at least α. The crisp set constructed in this way is called the α-cut of the fuzzy set *A*, and denoted as A_α.

Just as a fuzzy set extends a classical or crisp set, a fuzzy relation then extends the concept of a crisp relation from *X* to *Y*. For every (*x*,*y*) in *X* x *Y*, the quantity *R*(*x*,*y*) is interpreted as the strength of the existing *R*-link between *x* and *y*. In this way, a fuzzy strict preference (or indifference or incomparability) relation expresses the strength of a strict preference (or indifference or incomparability) among any couple of alternatives [8].

2.2. Analyzing a fuzzy quasi-order relation

Recall that a binary (fuzzy) relation *R* in a universe X is called a (fuzzy) *quasi-order relation* in X if and only if it is reflexive and transitive. The following characterization of a fuzzy quasi-order relation in terms of its α-cuts will prove to be important: *R* is a fuzzy quasi-order relation in X if and only if for all values of α (with α belonging to the interval [0,1]) it holds that R_α is a (crisp) quasi-order relation in X. Another crucial result is that any binary fuzzy relation *R* has a fuzzy quasi-order closure *Q*, i.e., a least inclusive fuzzy relation *Q* containing *R* and possessing the reflexivity and transitivity properties [1]. These results on fuzzy quasi-order relations allow us to reveal important information contained in the offers of the various buyers.

The starting point of our analysis are the preferences of the seller with respect to the various values specified in the buyers' offers for the different criteria of the product for sale. We assume that these preferences are given in the [0,1] interval, with 0 indicating no preference and 1 indicating a crisp preference, and conveniently summarize these preferences in a fuzzy preference relation *P*.

Next, we compare all rows in *P* two by two and compute how much every row is included in any other row. The degree to which one row is included in another row reflects how the seller's preferences on one buyer's values compared to her preferences on another buyer's values. We now summarize these degrees of inclusion in a binary fuzzy relation *D*, a relation in the set of buyers *B*.

More formally, and denoting the i-th row of *P* as the set P_i, we define the degree of inclusion of the i-th and j-th row of P as

$$Inc(P_i, P_j) = \frac{1}{n}\sum_k \min(1, 1 - P_{ik} + P_{jk}),$$

with $k=1,\ldots n$ the number of criteria, and P_{ik} the seller's degree of preference with respect to the value offered by buyer *i* on criterion *k*. As *D* is not necessarily transitive, we construct the fuzzy quasi-order closure Q of D in order to guarantee transitivity.

Since Q is a fuzzy quasi-order relation, we know that every α-cut of Q is a crisp quasi-order relation. The α-cuts of Q have the following interpretation: for any two buyers b_i and b_j, $(b_i, b_j) \in Q_\alpha$ if and only if b_i*'s offer is at most as good as* b_j's offer, with degree of confidence α.

Each α-cut of Q is a quasi-order relation in the set of buyers B. To Q_α corresponds an equivalence relation E_α in B defined by

$$(b_1, b_2) \in E_\alpha \Leftrightarrow (b_1, b_2) \in Q_\alpha \wedge (b_2, b_1) \in Q_\alpha.$$

The equivalence relation partitions the set of buyers into classes of buyers that are equally good. The equivalence class $[b_i]_\alpha$ of a buyer b_i is given by

$$[b_i]_\alpha = \{b_j \mid (b_i, b_j) \in E_\alpha\}.$$

The corresponding quotient set B_α is then given by $B_\alpha = \{[b]_\alpha \mid b \in B\}$. The quasi-order relation Q_α induces an order relation $\leq_\alpha$ in the quotient set, in the following way:

$$[b_i]_\alpha \leq_\alpha [b_j]_\alpha \Leftrightarrow (b_i, b_j) \in Q_\alpha.$$

Finally, we will make use of Hasse diagrams to graphically represent the preferences, indifferences or in-comparabilities among the buyers defined by the order relation $\leq_\alpha$ at every level α [3].

3. Numerical illustration

3.1. Analysis of the buyers' offers

To illustrate the above theoretical analyis, we consider the following multi-criteria negotiation. A company's HR manager has posted a job opening on an electronic marketplace, and she has described that job in terms of three criteria C_1 (*Hourly Salary*), C_2 (*Duration*) and C_3 (*Starting date*). Moreover, she (or her company) has particular preferences on these criteria, for instance '10 US dollar', '2 weeks', and 'Next week', respectively. Shortly after posting the job, five people respond to the posted job offer. Every applicant offers particular values on the criteria, and these differ more or less significantly from the manager's own values. The magnitude of these value differences will determine the preferences the HR manager has with respect to these offers.

Assume that she has the following preferences, summarized in the preference relation P:

$$P = \begin{pmatrix} 0.1 & 0.3 & 0.2 \\ 0.2 & 0.5 & 0.1 \\ 0.7 & 0.8 & 0.3 \\ 0.1 & 0.7 & 0.5 \\ 0.6 & 0.5 & 0.5 \end{pmatrix},$$

where P(i,j) is the degree to which the HR manager prefers her own value on the j-th criterion to the value suggested in the counteroffer of the i-th buyer on the j-th criterion. For example, the value of 0.1 for P(1,1) indicates that she slightly prefers (to a degree 0.1) her own value of 10 dollar on 'Hourly Salary' to the Salary value proposed in the first applicant's counteroffer.

We should make two important observations here. First, it happens that in our particular example all values in P are strictly larger than zero. This means that none of the applicants has offered a value that is better for the seller, for instance working for 9 dollar per hour, or starting the next day. Secondly, the smaller a degree of preference in P is, the better the corresponding value of that offer is to the seller. This implies that the interpretation of the α-cuts of Q will be different : for any two applicants b_i and b_j, $(b_i, b_j) \in Q_\alpha$ if and only if b_i*'s offer is at least as good as* b_j's offer, with degree of confidence α.

The calculation of the inclusion degrees of the various rows P_i is straightforward, and leads to the following dependency relation D, a binary relation in the set of applicants:

$$D = \begin{pmatrix} 1 & 0.97 & 1 & 1 & 1 \\ 0.9 & 1 & 1 & 0.97 & 1 \\ 0.6 & 0.67 & 1 & 0.77 & 0.87 \\ 0.77 & 0.8 & 0.93 & 1 & 0.93 \\ 0.67 & 0.73 & 0.93 & 0.83 & 1 \end{pmatrix}.$$

The transitive closure Q of D then is given by:

$$Q = \begin{pmatrix} 1 & 0.97 & 1 & 1 & 1 \\ 0.9 & 1 & 1 & 0.97 & 1 \\ 0.8 & 0.8 & 1 & 0.83 & 0.87 \\ 0.8 & 0.8 & 0.93 & 1 & 0.93 \\ 0.8 & 0.8 & 0.93 & 0.83 & 1 \end{pmatrix}.$$

Note that we have quite a large difference for some positions in Q as compared to D – e.g. the most extreme change is for $Q(3,1)$: a change from 0.6 to 0.8. Indeed, in general, the transitive closure will strenghten the weakest links, i.e., increase the degrees of the weakest inclusions.

From Q, it follows that we must consider α-cuts at the following levels: {0.8,0.83,0.87,0.9,0.93,0.97,1}, leading to the following crisp cut relations:

$$Q_{0.83} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{pmatrix}, \quad Q_{0.87} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix},$$

$$Q_{0.9} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix}, \quad Q_{0.93} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix},$$

$$Q_{0.97} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad Q_{1} = \begin{pmatrix} 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Note that $Q_{0.8}$ is the identity matrix.

Let us now analyze Q at the given cut-levels. At α = 0.8, the HR manager is indifferent to all offers. This means that, at this level, no significant differences exist among the various offers she has received. At α = 0.83, we observe a first separation among the offers. The HR manager is still indifferent to Offers 1 and 2, as well as to Offers 3, 4 and 5. However, at this level, she prefers Offers 1 and 2 to Offers 3, 4 and 5. At α = 0.87, Offer 4 becomes more preferred than Offers 3 and 5, while remaining less preferred than Offers 1 and 2. At α = 0.90, the original group of indifferent Offers 3, 4 and 5 completely dissolves: Offer 4 is preferred to Offer 5 and Offer 5 is preferred to Offer 3, while all three remain less preferred than Offers 1 and 2. At α = 0.93, Offers 1 and 2 are now ‘separating’ – Offer 1 is preferred to Offer 2, and we obtain a linear ordering of all offers. At α = 0.97, the first incomparabilities show up: Offers 4 and 5 become incomparable, and Offer 3 becomes incomparable to both Offer 4 and Offer 5. Offer 1 is still the preferred Offer, while Offer 2 remains in second place. Finally, at the highest cut level α = 1, also Offer 1 and Offer 2 become incomparable, as well as Offer 2 and Offer 4, which means that Offer 2 separates from 1 and 4. We illustrate the various relationships among the offers at the different cuts in Figure 1.

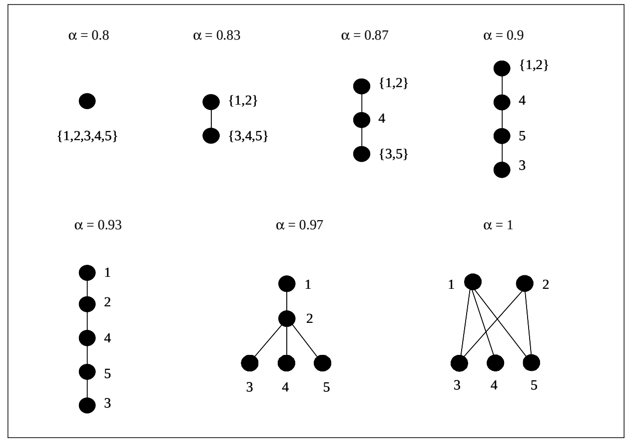

Figure 1. Hasse diagrams for Q_α, with $\alpha \in$ {0.8,0.83,0.87,0.9,0.93,0.97,1}.

3.2. Analysis of the seller’s criteria

Interestingly, we can analyze the seller’s criteria with the same technique. This analysis will show us on which criteria the seller’s values differ most from the values offered by the buyers. As such, the seller will be aware of which criteria are most discriminating among the various buyers, and on which criteria to focus for further negotiation. This type of analysis is essential when preparing to make trade-offs in the negotiation [5].

We obtain a ranking of the seller’s criteria by considering the transpose P^t of the original matrix P, i.e.,

$$P^t = \begin{pmatrix} 0.1 & 0.2 & 0.7 & 0.1 & 0.6 \\ 0.3 & 0.5 & 0.8 & 0.7 & 0.5 \\ 0.2 & 0.1 & 0.3 & 0.5 & 0.5 \end{pmatrix}.$$

Note that the smaller the preferences are on a criterion (a row in P^t), the better the buyers’ responses are to the seller on that particular criterion. The dependency matrix D derived from P^t is given by:

$$D = \begin{pmatrix} 1 & 0.98 & 0.88 \\ 0.76 & 1 & 0.76 \\ 0.9 & 1 & 1 \end{pmatrix},$$

which is a transitive relation, i.e., $Q = D$ in this case. The interpretation of the the α-cuts of Q in this case is: for any two criteria c_i and c_j, $(c_i, c_j) \in Q_\alpha$ if and only if c_i *is at least as good as* c_j, with degree of confidence α.

We now have to consider cuts at the following levels: {0.76, 0.88, 0.9, 0.98,1}. At α = 0.76, the lowest cut level, all alternatives (the criteria in this case) are indifferent. At α = 0.88, we have that Criterion 2 becomes worse than Criteria 1 and 3, both remaining indifferent. This means that Criterion 2 is the criterion on which worse scores are received from the buyers. At α = 0.9, we obtain a linear rank order: Criterion 3 is better than Criterion 1, which is better than Criterion 2. At α = 0.98, Criteria 1 and 3 have become incomparable, while Criterion 2 remains the worst criterion. Finally, at the highest cut level of α = 1, Criterion 3 is better than Criterion 2, but both are incomparable to 1.

From this analysis we can conclude that Criterion 2 overall is the worst criterion, meaning that the seller has the lowest preferences for the offers the buyers have made on this criterion. Criterion 3 is always better than Criterion 2, while the position of Criterion 1 varies.

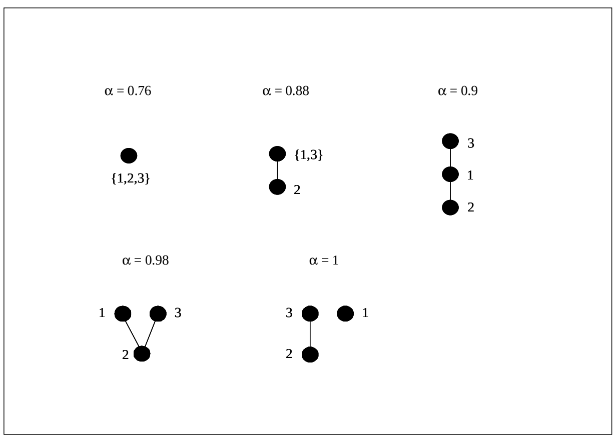

Figure 2. Hasse diagrams for Q_α, with $\alpha \in \{0.76, 0.88, 0.9, 0.98, 1\}$.

4. Conclusion

We have presented a model of individual decision making that better utilizes the information context of decision maker than the classic models. In this model the set of alternatives are evaluated through a combination of crisp and fuzzy relational analysis methods. Central to this model is the incomparability relation that occurs when the decision maker has conflicting information to make decisions. Here we have shown how the decision maker can, through varying the levels of the alpha cut, form different rank orderings of alternatives.

Our future goal is to append this individual decision making model with a conflict resolution / coordination protocol (such as negotiation or persuasion mechanism) [9]. Therefore, the original individual decision making problem is transformed from an interactive into a distributed one where individual decision models (strategically) negotiate with one another over the set of possible alternatives. One possible object of negotiation is the level of the alpha cut. Alternatively, information revealed in the course of negotiation can be used to determine the value of alpha. Scalable algorithms and architectures will be developed that implement this model for computational agents that operate autonomously as decision makers or semi-autonomously as decision support systems. Finally, the developed model and its multi-agent extension will be evaluated in bi-lateral bargaining and multi-lateral RFQ as well as other trading protocols.

Acknowledgements

The first author work is supported, in part, by the New Jersey Institute of Technology SBR2000 Faculty Grant.

References

[1] W. Bandler and L. Kohout, “Fuzzy relational products as a tool for analysis and synthesis of the behavior of complex natural and artificial systems”, in *Fuzzy Sets: Theory and Application to Policy Analysis and Information Systems* (Wang, S.K. and Wang, P.P., eds.), Plenum Press, New York and London 1980, pp. 341-367.

[2] M. Bichler, “An experimental analysis of multi-attribute auctions”, *Decision Support Systems* 29 (3) (2000), 249-268.

[3] B. De Baets and B. Van de Walle, “Dependencies among alternatives and criteria in MCDM problems”, *Proceedings of the Fourth International Workshop on Current Issues in Fuzzy Technologies* (1994), 80 - 85.

[4] J. Fodor and M. Roubens, *Fuzzy Preference Modeling and Multicriteria Decision Support*, Kluwer, Dordrecht, 1994.

[5] P. Faratin, C. Sierra and N. Jennings, “Using Similarity Criteria to Make Negotiation Trade-Offs”, *Proceedings of the International Conference on Multiagent Systems* (ICMAS-2000, Boston, MA), 2000, 119-126.

[6] H. Raiffa, *The art and science of negotiation*, Harvard University Press, 1982, Cambridge USA.

[7] C. Sierra, P. Faratin and N.R. Jennings, “Deliberative automated negotiators using fuzzy similarities”, in *Proceedings of EUSFLAT, Joint Conference on Fuzzy Logic*, 2000, 155-158.

[8] B. Van de Walle, B. De Baets and E. E. Kerre, “Characterizable Fuzzy Preference structures”, *Annals of Operations Research* (special issue on Preference Modelling) **80** (1998), 105 – 136.

[9] B. Van de Walle and P. Faratin, “Fuzzy preferences for multi-criteria negotiation”, position paper accepted for *AAAI Fall 2001 Symposium*, Massachussets, USA.