

INTERDISCIPLINARY APPROACH TO AUTOMATED NEGOTIATION: A PRELIMINARY REPORT

Fernando Lopes, A. Q. Novais
INETI, Dep. de Modelação e Simulação
Est. do Paço do Lumiar, 1649-038 Lisboa, Portugal
{fernando.lopes, augusto.novais}@ineti.pt

Helder Coelho
Universidade de Lisboa, Dep. de Informática
Campo Grande, 1749-016 Lisboa, Portugal
hcoelho@di.fc.ul.pt

KEYWORDS

Multi-agent systems, autonomous agents, automated negotiation, game theory, bargaining, efficiency.

ABSTRACT

Autonomous agents with negotiation competence are becoming increasingly important and pervasive. This paper follows an interdisciplinary approach to build autonomous negotiating agents by considering both game-theoretic techniques and bargaining procedures from the social sciences. The paper presents a generic model that handles bilateral multi-issue negotiation, describes equilibrium strategies for the bargaining game of alternating offers, and formalizes important strategies used by human negotiators. Autonomous agents equipped with the model are able to negotiate under both complete and incomplete information, thereby making them very compelling for automated negotiation.

INTRODUCTION

Automated negotiation systems with computational agents representing individuals or organizations and capable of reaching mutually beneficial agreements are becoming increasingly important and pervasive. Examples, to mention a few, include the business trend toward agent-based supply chain management, the pivotal role that electronic commerce is increasingly assuming in many organizations, and the industrial trend toward virtual enterprises.

Artificial intelligence (AI) researchers have paid some attention to automated negotiation over the last years and a number of models have been proposed in the literature. These models can be classified into three main classes (Jennings et al. 01):

- *game-theoretic models* – provide clear analysis of specific negotiation situations and precise results concerning the optimal strategies negotiators should choose, *i.e.*, the strategies that maximize the negotiation outcome (see, e.g., Kraus 01, Fatima et al. 06);

- *heuristic models* – provide general guidelines to assist negotiators and beneficial strategies for moving toward agreement, *i.e.*, strategies that lead to good (rather than optimal) outcomes of negotiation (see, e.g., Lopes et al. 04, Li et al. 06);
- *argumentation-based models* – allow negotiators to argue about their mental attitudes during the negotiation process. Thus, in addition to submitting proposals, negotiators can provide arguments either to justify their negotiation stance or to persuade other negotiators to change their negotiation stance (see, e.g., Rahwan et al. 04).

Game-theoretic models have some highly desirable properties, such as Pareto efficiency and the ability to guarantee convergence. However, most models make the following restrictive assumptions: (i) the agents are rational, (ii) the set of candidate solutions is fixed and known by all the agents, and (iii) each agent knows either the other agents' payoffs for all candidate solutions or the other agents' potential attitudes toward risk and expected-utility calculations.

Heuristic models are typically based on informal models of interaction and negotiation from the social sciences. They exhibit the following desirable features: (i) are based on realistic assumptions, and (ii) make use of moderate computational resources to find acceptable solutions. However, most models have a number of limitations. Firstly, they lack a rigorous theoretical underpinning – they are essentially ad hoc in nature. Secondly, they often lead to outcomes that are sub-optimal. Finally, they need extensive evaluation.

Argumentation-based models attempt mainly to marry the exchange of offers with the exchange of arguments. This permits great flexibility since, for instance, it makes possible to persuade agents to change their view of an offer during the course of negotiation. However, most models make considerable demands on any implementation, mainly because they appeal to very rich representations of the agents and their environments.

Against this background, this paper follows an interdisciplinary approach to build agents with negotiation competence. Its main purpose is to present a generic model of negotiation for autonomous agents. The model handles bilateral multi-issue negotiation, describes equilibrium strategies for the bargaining game of alternating offers, and formalizes a set of negotiation strategies frequently used by human negotiators.

Thus, this paper uses both game-theoretic techniques and negotiation procedures from the social sciences as a basis to develop autonomous negotiating agents. On the one hand, it considers two fully informed agents about the various aspects of the bargaining game of alternating offers and employs game-theoretic techniques to define equilibrium strategies. On the other hand, it considers two incompletely informed agents and formalizes important strategies studied in the social sciences (a fruitful area for future work will be to evaluate the effectiveness of these strategies in a number of different situations).

The remainder of the paper is structured as follows. The next section presents the key features of a generic model of negotiation for autonomous agents. In particular, we describe the negotiation protocol and the preferences of the negotiators, present equilibrium strategies for the bilateral multi-issue bargaining game of alternating offers, and introduce a number of concession strategies and their associated tactics. Following this, we discuss related work, present concluding remarks, and outline a number of issues which require further investigation.

THE NEGOTIATION MODEL

Negotiation is usually understood as proceeding through three distinct phases (Lewicki et al. 03): a beginning or initiation phase, a middle or problem-solving phase, and an ending or resolution phase. The initiation phase focuses on preparation and planning for negotiation and is marked by each party's efforts to posture for positions. The problem-solving phase seeks a solution for the dispute and is characterized by movement toward a final agreement. The resolution phase focuses on implementing the final agreement.

Effective preparation and planning involves the creation of a well-laid plan specifying the activities that negotiators should attend to before starting to negotiate. That plan, and the confidence derived from it, is often a critical factor for achieving negotiation objectives. Accordingly, we describe below various activities that negotiators make efforts to perform in order to carefully prepare and plan for negotiation (see Lopes et al. 02, for an in-depth discussion).

Let $Ag = \{ag_1, ag_2\}$ be the set of autonomous negotiating agents. Let $Agenda = \{is_1, \dots, is_n\}$ be the negotiating agenda – the final set of issues to be deliberated. Effective planning requires that negotiators prioritize the issues. In this work, priorities are set by ranking-order the issues, *i.e.*, by defining the most important, the second most important, and so on. The priority pr_{il} of an agent $ag_i \in Ag$ for each issue $is_l \in Agenda$ is a number that represents its order of preference. The weight w_{il} of is_l is a number that represents its relative importance.

Effective planning also requires that negotiators define two key points for each issue at stake: the limit or reservation price and the target point or level of aspiration. The limit is the point where a negotiator decides that it should stop to negotiate, because any settlement beyond this point is not minimally acceptable. It is set by determining realistic alternatives to a negotiated agreement and selecting the best alternative. The target point is the point where a negotiator realistically expects to achieve a settlement. It is set by first defining the limit and then brainstorming about what can be realistically possible. The limit of ag_i for is_l is denoted by lim_{il} and the target point by trg_{il} .

Additionally, effective planning requires that negotiators agree on an appropriate protocol that defines the rules governing the interaction. The protocol can be simple, allowing agents to exchange only proposals. Alternatively, the protocol can be sophisticated, allowing agents to provide arguments to support their negotiation stance. As noted, most sophisticated protocols make considerable demands on any implementation. Thus, in this work we consider a simple protocol (see subsection “The Negotiation Protocol”).

Finally, effective planning requires that negotiators be able to select appropriate strategies that account for their individual actions. Negotiation strategies can reflect a variety of behaviours and lead to strikingly different outcomes. Some strategies are in equilibrium, meaning that no designer will benefit by building agents that use any other strategies when it is known that some agents are using equilibrium strategies. Thus, for some situations of complete information, the agents can be designed to adopt equilibrium strategies (see subsection “Equilibrium Strategies”). However, for situations of incomplete information, the problem of strategic choice is rather complex. In these situations, many bargaining models have different equilibria sustained by different assumptions on what an individual in the game would believe if its opponent took an action that it was not supposed to take in equilibrium. Hence, our study differs from this line of work. As stated, we address the challenge of building

agents that are able to negotiate under incomplete information by formalizing relevant strategies used by human negotiators and empirically evaluating the effectiveness of these strategies in different situations (see subsection “Concession Strategies and Tactics”).

At this stage, we hasten to add an explanatory and cautionary note. This paper primarily focuses on the problem-solving phase of negotiation. Most AI researchers have also focused on this phase of negotiation and have addressed various issues associated with the design of negotiation protocols and strategies. Their work has resulted in some advances in automated negotiation, but it is clear that much more research still needs to be performed. At present, the deployment of agents that are able to negotiate under both complete and incomplete information is still in its infancy.

The Negotiation Protocol

The protocol is an alternating offers protocol (Osborne and Rubinstein 90). Two agents or players, ag_1 and ag_2 , bargain over the division of the surplus of $n \geq 2$ issues (or goods). The players determine an allocation of the issues by alternately proposing offers at times in $\mathcal{T} = \{1, 2, \dots\}$. This means that one offer is made per time period $t \in \mathcal{T}$, with an agent, say ag_i , offering in odd periods $\{1, 3, \dots\}$, and the other agent ag_j offering in even periods $\{2, 4, \dots\}$.

The negotiation procedure, labelled the “joint-offer procedure”, involves bargaining over the allocation of the entire endowment stream at once. An offer is a vector (x_1, \dots, x_n) specifying a division of the n goods. Once an agreement is reached, the agreed-upon allocations of the goods are implemented. This procedure permits agents to exploit the benefits of trading-off concessions on their less preferred goods for concessions by their opponent on the more preferred goods.

The players’ preferences are modelled by assuming that each player ag_i discounts future payoffs at some given rate δ_i^t , $0 < \delta_i^t < 1$, (δ_i^t is referred to as the discount factor and the preferences as time preferences with a constant discount rate). The cost of bargaining derives from the delay in consumption implied by a rejection of an offer. Practically speaking, the justification for this form of preferences takes into account the fact that money today can be used to make money tomorrow.

Let U_i be the payoff function of ag_i . For simplicity and tractability, we assume that U_i is separable in all their arguments and that the per-period delay costs are the same for all issues:

$$U_i(is_{i1}, \dots, is_{in}, t) = \delta_i^{(t-1)} \sum_{l=1}^n w_{il} u_{il}(is_{il})$$

where w_{il} is the weight of an issue is_{il} and u_{il} is the component payoff function for is_{il} , *i.e.*, the function that gives the payoff that ag_i assigns to a value of is_{il} . The function u_{il} is a continuous, strictly monotonic, and linear function. The distinguish feature of time preferences with a constant discount rate is the linearity of the function u_{il} . The payoff of disagreement is normalized at 0 for both players.

Equilibrium Strategies

Game theory can provide sound design principles for computer scientists. Consider two fully informed players bargaining over four distinct goods $\{X_1, \dots, X_4\}$. Each good is modelled as an interval $[0, 1]$ (or as a divisible pie of size 1). The players’ preferences are as follows:

$$U_i = \delta_i^{(t-1)} (a x_1 + b x_2 + x_3 + x_4)$$

$$U_j = \delta_j^{(t-1)} [(1-x_1) + (1-x_2) + c(1-x_3) + d(1-x_4)]$$

where x_l , $l=1, \dots, 4$, denotes the share of ag_i for each pie. The parameters a , b , c , and d allow the marginal utilities of the players to differ across issues and players. We consider $a > b > 1$ and $d > c > 1$, *i.e.*, ag_i places greater emphasis on goods X_1 and X_2 while ag_j values goods X_3 and X_4 more. Also, we consider that δ_i and δ_j are close to 1 and the parameters a , b , c , and d are close to one another. Let $p_{j \rightarrow i}^{t-1}$ and $p_{i \rightarrow j}^t$ denote the offers that ag_j proposes to ag_i in period $t-1$ and ag_i proposes to ag_j in period t , respectively. Consider the following strategies for ag_i and ag_j :

$$str_i^* = \begin{cases} \text{offer } (1, 1, x_{i3}^*, 0) & \text{if } ag_i \text{'s turn} \\ \text{if } U_i(p_{j \rightarrow i}^{t-1}) \geq U_i^* \text{ accept else reject} & \text{if } ag_j \text{'s turn} \end{cases}$$

$$str_j^* = \begin{cases} \text{offer } (1, x_{j2}^*, 0, 0) & \text{if } ag_j \text{'s turn} \\ \text{if } U_j(p_{i \rightarrow j}^t) \geq U_j^* \text{ accept else reject} & \text{if } ag_i \text{'s turn} \end{cases}$$

where $U_i^* = U_i(1, x_{j2}^*, 0, 0)$, $U_j^* = U_j(1, 1, x_{i3}^*, 0)$, and the shares are the following:
 $x_{i3}^* = \frac{\delta_i \delta_j (a+b) - \delta_j (a+b+bc+bd) + bc + bd}{bc - \delta_i \delta_j}$ and $x_{j2}^* = \frac{\delta_i (\delta_i \delta_j (a+b) - \delta_j (a+b+bc+bd) + bc + bd) + (bc - \delta_i \delta_j) (a \delta_i + b \delta_i - a)}{b(bc - \delta_i \delta_j)}$.

Remark 1. For the two-sided four-issue bargaining game of alternating offers with an infinite horizon, in which the players’ preferences are as described above, the pair of strategies (str_i^*, str_j^*) form an equilibrium. The outcome is the following:

$$x_1^* = 1, \quad x_2^* = 1, \quad x_4^* = 0,$$

$$x_3^* = \frac{\delta_i \delta_j (a+b) - \delta_j (a+b+bc+bd) + bc + bd}{bc - \delta_i \delta_j}$$

Agreement is immediately reached with no delay. The outcome is Pareto optimal. Letting $\delta_i \rightarrow 1$ and $\delta_j \rightarrow 1$, the equilibrium division is $(1, 1, 0, 0)$.

The formal proof is based on the familiar necessary conditions for equilibrium: ag_i is indifferent between waiting one period to have its offer accepted and accepting ag_j 's offer immediately, and ag_j is indifferent between waiting one period to have its offer accepted and accepting ag_i 's offer immediately. Let $\mathbf{x}_i^* = (x_{i1}^*, \dots, x_{i4}^*)$ and $\mathbf{x}_j^* = (x_{j1}^*, \dots, x_{j4}^*)$ be the equilibrium proposals of ag_i and ag_j , respectively. The problem for ag_i is to find an offer that maximizes its payoff (because it is a payoff maximizer) subject to being acceptable to its opponent, *i.e.*,

maximize:

$$U_i(x_1, \dots, x_4, t) = \delta_i^{(t-1)} (ax_1 + bx_2 + x_3 + x_4)$$

subject to:

$$\begin{aligned} (1-x_{i1}^*) + (1-x_{i2}^*) + c(1-x_{i3}^*) + d(1-x_{i4}^*) = \\ \delta_j[(1-x_{j1}^*) + (1-x_{j2}^*) + c(1-x_{j3}^*) + d(1-x_{j4}^*)] \\ 0 \leq x_{il}^* \leq 1, \quad 0 \leq x_{jl}^* \leq 1, \quad \text{for } l=1, \dots, 4 \end{aligned}$$

The problem for ag_j is stated in a similar way and is omitted. Solving both maximization problems yields the outcome specified in the statement of the Remark. In the limit, letting $\delta_i \rightarrow 1$ and $\delta_j \rightarrow 1$, the outcome of the equilibrium is $(1, 1, 0, 0)$. This outcome is on the Pareto frontier and corresponds to the utility pair $(a+b, c+d)$.

Now, consider two fully informed players bargaining over n distinct goods $\{X_1, \dots, X_n\}$. Again, each good is modelled as a divisible pie of size 1. We consider that the players set different weights for the goods such that: $w_{i1}/w_{j1} > w_{i2}/w_{j2} > \dots > w_{in}/w_{jn}$.

Remark 2. The bilateral multi-issue bargaining game of alternating offers with an infinite horizon, in which the players' preferences are as described above, has an equilibrium. The outcome is Pareto optimal:

$$(x_1^*, \dots, x_{k-1}^*, x_k^*, x_{k+1}^*, \dots, x_n^*) = (1, \dots, 1, s, 0, \dots, 0)$$

where x_l^* , $l=1, \dots, n$, denotes the share of ag_i for each divisible pie. The constant s represents the share of ag_i for the X_k pie.

The formal proof is also based on the familiar necessary conditions for equilibrium. In short, consider that $\mathbf{x}_i^* = (x_{i1}^*, \dots, x_{in}^*)$ and $\mathbf{x}_j^* = (x_{j1}^*, \dots, x_{jn}^*)$ represent the equilibrium proposals of players ag_i and ag_j , respectively. The problem for ag_i is stated as follows:

maximize:

$$U_i(x_1, \dots, x_n, t) = \delta_i^{(t-1)} \sum_{l=1}^n w_{il} x_l$$

subject to:

$$\begin{aligned} U_j(x_{j1}^*, \dots, x_{jn}^*, t) = U_j(x_{j1}^*, \dots, x_{jn}^*, t+1) \\ 0 \leq x_{il}^* \leq 1, \quad 0 \leq x_{jl}^* \leq 1, \quad \text{for } l=1, \dots, n \end{aligned}$$

This maximization problem is similar to the continuous knapsack problem and solvable by a greedy approach (see, e.g., Martello and Toth 90). First, ag_i gives away the maximum possible share of the issue with the lowest ratio of weights. If the supply of that issue is exhausted, it gives away the maximum possible share of the issue with the next lowest ratio of weights, and so forth until ag_j gets the utility of $U_j(x_{j1}^*, \dots, x_{jn}^*, t+1)$. The problem for ag_j is also analogous to the continuous knapsack problem and solvable by a greedy approach.

Concession Strategies and Tactics

Human negotiation practice can provide rules-of-thumb to agent designers. Consider two incompletely informed agents bargaining over n distinct issues $\{is_1, \dots, is_n\}$. For convenience, each issue is_l is modelled as an interval $[min_l, max_l]$. The agents' preferences are as defined in the subsection "The Negotiation Protocol".

The opening stance (or position) and the pattern of concessions are two central elements of real-world negotiation. Three different opening positions (extreme, reasonable and modest) and three levels of concession magnitude (large, moderate and small) are commonly discussed in the literature (Lewicki et al. 03, Pruitt and Kim 04). They can lead to a number of concession strategies, notably:

- *starting high and conceding slowly* – negotiators adopt an optimistic opening attitude and make successive small concessions;
- *starting reasonable and conceding moderately* – negotiators adopt a realistic opening attitude and make successive moderate concessions.

These strategies are implemented through a variety of tactics. The terms "strategy" and "tactic" differ in scope – a tactic is a short-term move designed to enact a high-level strategy.

Let $p_{j \rightarrow i}^{t_{n-1}}$ be the offer that ag_j has proposed to ag_i in period t_{n-1} . Likewise, let $p_{i \rightarrow j}^{t_n}$ be the offer that ag_i is ready to propose in the next time period t_n . The formal definition of a generic concession strategy follows.

Definition 1. Let $ag_i \in Ag$ be a negotiating agent. A concession strategy for ag_i is a function that specifies either the tactic to apply at the beginning of negotiation or the tactic that defines the concessions to be made during the course of negotiation:

$$s = \begin{cases} \text{apply } tact_i^{t_1} & \text{if } ag_i \text{'s turn and } t=1 \\ \text{apply } tact_i^t & \text{if } ag_i \text{'s turn and } t > 1 \\ \text{if } U_i(p_{j \rightarrow i}^{t_{n-1}}) \geq U_i(p_{i \rightarrow j}^{t_n}) \text{ accept else reject if } ag_j \text{'s turn} \end{cases}$$

where $tact_i^{t_1}$ is an opening negotiation tactic and $tact_i^t$ is a concession tactic. ■

The two concession strategies are defined by considering different tactics. For instance, the “starting reasonable and conceding moderately” strategy is defined by: “ $tact_i^{t_1} = starting_realistic$ ” and “ $tact_i^t = moderate$ ”.

Opening negotiation tactics are functions that specify the initial values for each issue is_l at stake. The following three tactics are commonly used in real-life negotiations (Lewicki et al. 03):

- *starting optimistic* – specifies a value far from the target point;
- *starting realistic* – specifies a value close to the target point;
- *starting pessimistic* – specifies a value close to the limit.

The definition of the tactic “starting realistic” follows (the definition of the other two tactics is essentially identical, and is omitted).

Definition 2. Let $ag_i \in Ag$ be a negotiating agent and $is_l \in Agenda$ a negotiation issue. Let trg_{il} be the target point of ag_i for is_l . The tactic starting realistic for ag_i is a function that takes is_l and trg_{il} as input and returns the initial value $v[is_l]_i^{t_1}$ of is_l :

$$starting_realistic(is_l, trg_{il}) = v[is_l]_i^{t_1}$$

where $v[is_l]_i^{t_1} \in [trg_{il} - \epsilon, trg_{il} + \epsilon]$ and $\epsilon > 0$ is small. ■

Concession tactics are functions that compute new values for each issue is_l . The following five tactics are commonly used (Lewicki et al. 03):

- *stalemate* – models a null concession on is_l ;
- *tough* – models a small concession on is_l ;
- *moderate* – models a moderate concession;
- *soft* – models a large concession;
- *accommodate* – models a complete concession.

The definition of a generic concession tactic follows (without loss of generality, we consider that ag_i wants to maximize is_l).

Definition 3. Let $ag_i \in Ag$ be a negotiating agent, $is_l \in Agenda$ a negotiation issue, and lim_{il} the limit of is_l . Let $v[is_l]_i^t$ be the value of is_l offered by ag_i at t . A concession tactic for ag_i is a function that takes $v[is_l]_i^t$, lim_{il} and the concession factor $Cf \in [0, 1]$ as input and returns the new value $v[is_l]_i^{t+2}$ of is_l :

$$concession_tactic(v[is_l]_i^t, lim_{il}, Cf) = v[is_l]_i^{t+2}$$

where $v[is_l]_i^{t+2} = v[is_l]_i^t - Cf(v[is_l]_i^t - lim_{il})$. ■

The five tactics are defined by considering different values for Cf . In particular, the stalemate tactic by $Cf=0$, the accommodate tactic by $Cf=1$, and the other three tactics by different ranges of values for Cf (e.g., the tough tactic by $Cf \in]0.00, 0.05]$, the moderate tactic by $Cf \in]0.05, 0.10]$, and the soft tactic by $Cf \in]0.10, 0.15]$).

RELATED WORK

AI researchers have developed a number of negotiation models incorporating specific protocols and libraries of negotiation strategies (see, e.g., Kraus 01, Lopes et al. 04, Fatima et al. 06, Li et al. 06). However, despite these and other relevant pieces of work, we are aware of no similar efforts to define strategies as functions that specify the tactics to be used at every period of negotiation. Tactics, in turn, are defined as functions that specify the short-term moves to be made throughout negotiation. Also, we are particularly interested in both equilibrium strategies that lead to Pareto optimal outcomes and important strategies studied in the social sciences and frequently used by human negotiators.

CONCLUSION

This article has followed an interdisciplinary approach and has used both game-theoretic techniques and negotiation procedures from the social sciences as a basis to develop autonomous negotiating agents. It has presented a model that handles bilateral multi-issue negotiation, describes equilibrium strategies for the bargaining game of alternating offers, and formalizes a set of human negotiation strategies. Our aim for the future is: (i) to extend the model, (ii) to empirically evaluate the model, and (iii) to use game-theoretic techniques as a basis to develop autonomous agents that are able to negotiate under incomplete information.

REFERENCES

- [Fatima et al. 06] Fatima S., M. Wooldridge and N. Jennings (2006), Multi-Issue Negotiation with Deadlines, *Journal of Artificial Intelligence Research*, 27:381–417.
- [Jennings et al. 01] Jennings N., P. Faratin, A. Lomuscio, S. Parsons, M. Wooldridge and C. Sierra (2001), Automated Negotiation: Prospects, Methods and Challenges, *Group Decision and Negotiation*, 10:199–215.
- [Kraus 01] S. Kraus (2001), *Strategic Negotiation in Multi-Agent Environments*, MIT Press.
- [Lewicki et al. 03] Lewicki R., B. Barry, D. Saunders and J. Minton (2003), *Negotiation*, McGraw Hill.
- [Li et al. 06] C. Li, J. Giampapa, and K. Sycara (2006), Bilateral Negotiation Decisions with Uncertain Dynamic Outside Options, *IEEE Transactions on Systems, Man, and Cybernetics (Part-C)*, 36:31–44.
- [Lopes et al. 02] F. Lopes, N. Mamede, A. Q. Novais, and H. Coelho (2002), A Negotiation Model for Autonomous Agents: Formal Description and Empirical Evaluation, *Journal of Intelligent & Fuzzy Systems*, 12:195–212.
- [Lopes et al. 04] Lopes F., N. Mamede, A. Q. Novais and H. Coelho (2004), Negotiation Strategies for Autonomous Computational Agents, In: *ECAI-04*, pp. 38–42.
- [Martello and Toth 90] Martello S. and P. Toth (1990), *Knapsack Problems: Algorithms and Computer Implementations*, John Wiley & Sons.
- [Osborne and Rubinstein 90] Osborne M. and A. Rubinstein (1990), *Bargaining and Markets*, Academic Press.
- [Pruitt and Kim 04] D. Pruitt and S. Kim (2004), *Social Conflict*, McGraw Hill.
- [Rahwan et al. 04] Rahwan I., S. Ramchurn, N. Jennings, P. McBurney, S. Parsons and L. Sonenberg (2004), Argumentation-based Negotiation, *The Knowledge Engineering Review*, 18:343–375.