

Towards an Interdisciplinary Framework for Automated Negotiation

Fernando Lopes¹, A.Q. Novais¹, and Helder Coelho²

¹ INETI, Dep. de Modelação e Simulação, Est. Paço Lumiar, 1649-038 Lisboa, Portugal
{fernando.lopez, augusto.novais}@ineti.pt

² Universidade de Lisboa, Dep. de Informática, Campo Grande, 1749-016 Lisboa, Portugal
hcoelho@di.fc.ul.pt

Abstract. Negotiation is an important and pervasive form of social interaction. The design of autonomous negotiating agents involves the consideration of insights from multiple relevant research areas to integrate different perspectives on negotiation. As a starting point for an interdisciplinary research effort, this paper presents a model that handles bilateral multi-issue negotiation, employs game-theoretic techniques to define equilibrium strategies for the bargaining game of alternating offers, and formalizes a set of negotiation strategies and tactics studied in the social sciences. Autonomous agents equipped with the model are currently being developed using the Jade framework. The agents are able to negotiate under both complete and incomplete information, thereby making the model in particular and the agents in general very compelling for automated negotiation.

1 Introduction

Autonomous negotiating agents representing individuals or organizations and capable of reaching mutually beneficial agreements are becoming increasingly important. 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 organizations, and the industrial trend toward virtual enterprises (see, e.g., [2,4]).

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 [7]: (i) game-theoretic models, (ii) heuristic models, and (iii) argumentation-based models. 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 negotiation outcome (see, e.g., [5,8]). 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 (see, e.g., [3,6]). Typically, they are based on informal models of interaction and negotiation from the social sciences. 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., [16]).

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 exhibit the following desirable features: (i) they are based on realistic assumptions, and (ii) they make use of moderate computational resources to find acceptable solutions (according to the principles of bounded rationality [18]). However, most models lack a rigorous theoretical underpinning – they are essentially ad hoc in nature. Also, they often lead to outcomes that are sub-optimal. Finally, there is often no precise understanding of how and why they behave the way they do. Consequently, 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. Accordingly, some researchers pointed out that agents which can argue in support of their negotiations will only ever represent a small, though important, class of automated negotiators [7].

Automated negotiation promises a higher level of process efficiency and a higher quality of agreements (when compared to traditional, face-to-face negotiation). In practice, the task of designing and implementing autonomous negotiating agents involves the consideration of insights from multiple relevant research areas to integrate different perspectives on negotiation [1]. Yet, most existing models primarily use either game-theoretic techniques or methods from the social sciences as a basis to develop negotiating agents, and largely ignore the integration of the results from both areas.

This paper argues that an interdisciplinary approach towards the development of autonomous negotiating agents is possible and highly desirable – game-theoretic (strategic) and behavioural negotiation theories can mutually reinforce each other and lead to more comprehensive and richer models. As a starting point for this research effort, this paper presents a model for autonomous agents that handles bilateral multi-issue negotiation, introduces equilibrium strategies for the bargaining game of alternating offers, and formalizes a set of strategies and tactics frequently used by human negotiators. On the one hand, it considers two fully informed agents about the various aspects of the bargaining game and employs game-theoretic techniques to define equilibrium strategies. On the other hand, it considers two incompletely informed agents about the various aspects of the game and formalizes a set of negotiation strategies and tactics studied in the social sciences.

This paper builds on our previous work in the area of negotiation [10,11,12]. In particular, it extends our negotiation model by introducing a number of negotiation strategies and tactics motivated by human procedures typical of multi-issue negotiation. It also lays the foundation for performing an experiment to investigate the performance of agents equipped with the model in terms of quality and cost of bargaining.

The remainder of the paper is structured as follows. Section 2 presents the negotiation model. In particular, subsection 2.1 describes the negotiation protocol and the preferences of the negotiators, subsection 2.2 introduces equilibrium strategies for the bilateral multi-issue bargaining game of alternating offers, and subsection 2.3 formalizes concession and problem solving strategies frequently used by human negotiators. Section 3 discusses related work and section 4 presents concluding remarks.

2 The Negotiation Model

Negotiation is usually understood as proceeding through three phases [9]: 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 [17,19]. 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 our earlier work for an in-depth discussion [10]).

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 set of (independent) issues at stake. The issues are quantitative variables, defined over continuous intervals. Effective planning requires that negotiators prioritize the issues, define the limits, and specify the targets. Priorities are set by rank-ordering 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. The limit lim_{il} is the point where ag_i decides that it should stop to negotiate, because any settlement beyond this point is not minimally acceptable. The level of aspiration or target point trg_{il} is the point where ag_i realistically expects to achieve a settlement.

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 earlier, most sophisticated protocols make considerable demands on any implementation. Thus, in this work we consider a simple protocol (see subsection 2.1).

Finally, effective planning requires that negotiators be able to select appropriate strategies that account for their individual actions. Traditionally, AI researchers have paid little attention to this pre-negotiation step. In the last several years, however, a number of researchers have developed models that include libraries of negotiation strategies (see, *e.g.*, [6,7,8]). 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 (see, e.g., [14] for an in-depth description of the standard game-theoretic concept of equilibrium). Thus, for some situations of complete information, the agents can be designed to adopt equilibrium strategies (see subsection 2.2). 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 – 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 2.3 and the comments on future work in section 4).

2.1 The Negotiation Protocol and Time Preferences

The negotiation protocol is an alternating offers protocol [14]. Two agents or players bargain over the division of the surplus of $n \geq 2$ issues (goods or pies). 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 concessions on different issues.

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:

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

where w_{il} is the weight of is_{il} and x_l denotes the share of ag_i for is_{il} . The component payoff function u_{il} for is_{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} [14]. The payoff of disagreement is normalized at 0 for both players.

2.2 Equilibrium Strategies

The negotiation process is modelled as an extensive game. For theoretical convenience, we consider the standard game-theoretic situation of two players completely informed about the various aspects of the game. The players are assumed to be rational, and each

player knows that the other acts rationally. Also, we consider settings in which there is more than one issue and the players have different evaluations of the issues.

First, consider a two-sided four-issue bargaining situation. Two players are jointly endowed with a single unit of each of four goods, $\{X_1, \dots, X_4\}$, and alternate proposals until they find an agreement. 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 and $(1 - x_l)$, $l=1, \dots, 4$, denote the shares of ag_i and ag_j for each pie, respectively. 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:

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

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

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:

$$x_1^* = 1, \quad x_2^* = 1, \quad x_3^* = \frac{\delta_i \delta_j (a+b) - \delta_j (a+b+bc+bd) + bc+bd}{bc - \delta_i \delta_j}, \quad x_4^* = 0$$

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 $(1, 1, 0, 0)$.

Now, consider a two-sided n -issue bargaining situation. Two players bargain over n distinct goods, $\{X_1, \dots, X_n\}$, and are allowed to alternate proposals until they find an agreement. Their preferences are as defined in subsection 2.1. Again, each good is modelled as a divisible pie of size 1. 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 proofs of Remark 1 and Remark 2 are 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. For instance, consider the n -issue bargaining situation. Let $\mathbf{x}_i^* = (x_{i1}^*, \dots, x_{in}^*)$ and $\mathbf{x}_j^* = (x_{j1}^*, \dots, x_{jn}^*)$ be the equilibrium proposals of ag_i and ag_j , respectively. The problem for ag_i is stated as follows:

$$\begin{aligned} \text{maximize:} \quad & U_i(x_1, \dots, x_n, t) = \delta_i^{(t-1)} \sum_{l=1}^n w_{il} x_l \\ \text{subject to:} \quad & U_j(x_{j1}^*, \dots, x_{jn}^*, t) = U_j(x_{j1}^*, \dots, x_{jn}^*, t+1) \\ & 0 \leq x_{il}^* \leq 1, \quad 0 \leq x_{jl}^* \leq 1, \quad \text{for } l=1, \dots, n \end{aligned}$$

This maximization problem is similar to the continuous (or fractional) knapsack problem and solvable by a greedy approach [13]. 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 stated in a similar way and also solvable by a greedy approach.

2.3 Concession and Problem Solving Strategies

Game theory can provide sound design principles for computer scientists. The last subsection has considered two fully informed agents and used game-theoretic techniques to define equilibrium strategies. The agents were creative, honest, and able to settle for the outcome that maximizes their benefit (resources were not wasted and money was not squandered). Yet, the assumption of complete information is of limited use to the designers of agents. In practice, agents have private information. Also, simple casual observation reveals the existence of concessions and long periods of disagreement in many actual negotiations. Furthermore, one agent, say ag_i , may wish to act rationally, but the other agent may not behave as a strategically sophisticated, utility maximizer — thus rendering conventional equilibrium analysis inapplicable.

Behavioral negotiation theory can provide rules-of-thumb to agent designers. The danger is that the designers may not be fully aware of the circumstances to which human practice is adapted, and hence use rules that can be badly exploited by new agents. Nevertheless, an increasing number of researchers consider that human practice is crucial to automated negotiation (see, e.g., [1]). There is a need to integrate the procedures and results from behavioral negotiation theory in bargaining models incorporating game-theoretic techniques. Accordingly, this subsection considers two incompletely informed agents about the various aspects of the bargaining game and formalizes relevant strategies studied in the social sciences.

Negotiation strategies can reflect a variety of behaviours and lead to strikingly different outcomes. However, the following two fundamental groups of strategies are commonly discussed in the behavioral negotiation literature [15,19]:

1. *concession making* – negotiators who employ strategies in this group reduce their aspirations to accommodate the opponent;
2. *problem solving* – negotiators maintain their aspirations and try to find ways of reconciling them with the aspirations of the opponent.

Although it is important to distinguish among these two groups of strategies, we hasten to add several explanatory notes. First, most strategies are implemented through a variety of tactics. The line between strategies and tactics often seems indistinct, but one major difference is that of scope. Tactics are short-term moves designed to enact or pursue broad (high-level) strategies [9]. Second, concession making strategies are essentially unilateral strategies – the decision to concede is fundamentally a unilateral one. By contrast, problem solving strategies are essentially social strategies. Third, most negotiation situations call forth a combination of strategies from different groups. Finally, most strategies are only informally discussed in the behavioral literature. They are not formalized, as typically happens in the game-theoretic literature.

Concession making behaviour aims at partially or totally accommodating the other party. 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 subsection 2.1. The opening stance and the pattern of concessions are two central elements of negotiation. Three different opening positions (extreme, reasonable and modest) and three levels of concession magnitude (large, moderate and small) are commonly discussed in the behavioral literature [9]. They can lead to a number of concession strategies, notably:

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

Let $p_{j \rightarrow i}^{t-1}$ be the offer that ag_j has proposed to ag_i in period $t-1$. Likewise, let $p_{i \rightarrow j}^t$ be the offer that ag_i is ready to propose in the next time period t . 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:*

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

where $tact_i^1$ is an opening negotiation tactic and $tact_i^t$ is a concession tactic.

The two aforementioned concession strategies are defined by considering different tactics. For instance, the “starting reasonable and conceding moderately” strategy is defined by: “ $tact_i^1 = \text{starting_realistic}$ ” and “ $tact_i^t = \text{moderate}$ ” (but see below).

Problem solving behaviour aims at finding agreements that appeal to all sides, both individually and collectively. The host of problem solving strategies includes [15]:

1. *low-priority concession making* – negotiators hold firm on more important issues while conceding on less important issues;
2. *logrolling* – negotiators agree to trade-off among the issues under consideration so that each party concedes on issues that are of low priority to itself and high priority to the other party.

Low-priority concession making involves primarily the analysis of one’s priorities and further concessions on less important issues. However, effective logrolling requires information about the two parties’ priorities so that concessions can be matched up. This information is not always easy to get. The main reason for this is that negotiators often try to conceal their priorities for fear that they will be forced to concede on issues of lesser importance to themselves without receiving any repayment [15]. Despite this, research evidence indicates that it is often not detrimental for negotiators to disclose information that can reveal their priorities – a simple rank order of the issues does not put negotiators at a strategic disadvantage [19]. Hence, we consider that negotiators willingly disclose information that can help to identify their priorities (e.g., their interests).

Logrolling can be viewed as a variant of low-priority concession making in which the parties’ priorities are in the opposite direction. The formal definition of a generic logrolling strategy follows (the definition of a low-priority concession making strategy is essentially identical, and is omitted).

Definition 2. Let $ag_i \in Ag$ be a negotiating agent and $ag_j \in Ag$ be its opponent. Let $Agenda$ denote the negotiating agenda, $Agenda^\oplus$ the subset of the agenda containing the issues of high priority for ag_i (and low priority for ag_j), and $Agenda^\ominus$ the subset of the agenda containing the issues of low priority for ag_i (and high priority for ag_j). A logrolling strategy for ag_i is a function that specifies either the tactic to apply at the beginning of negotiation or the tactics to make trade-offs during the course of negotiation:

$$\log \stackrel{def}{=} \begin{cases} \text{apply } tact_i^1 & \text{if } ag_i \text{'s turn and } t=1 \\ \text{apply } tact_i^{t\oplus} \text{ and } tact_i^{t\ominus} & \text{if } ag_i \text{'s turn and } t > 1 \\ \text{if } U_i(p_{j \rightarrow i}^{t-1}) \geq U_i(p_{i \rightarrow j}^t) \text{ accept else reject} & \text{if } ag_j \text{'s turn} \end{cases}$$

where $tact_i^1$ is an opening negotiation tactic, $tact_i^{t\oplus}$ is a concession tactic (to apply to the issues on $Agenda^\oplus$), and $tact_i^{t\ominus}$ is another concession tactic (to apply to the issues on $Agenda^\ominus$). ■

A number of logrolling strategies can be defined simply by considering different tactics. For instance, a strategy that specifies a realistic opening attitude, followed by null concessions on issues on $Agenda^\oplus$, and large concessions on issues on $Agenda^\ominus$, is defined by: “ $tact_i^1 = \textit{starting_realistic}$ ”, “ $tact_i^{t\oplus} = \textit{stalemate}$ ”, and “ $tact_i^{t\ominus} = \textit{soft}$ ” (but see below).

Opening negotiation tactics are functions that specify the initial values for each issue is_l at stake. The following three tactics are commonly discussed in the behavioral literature [9]:

1. *starting optimistic* – specifies a value far from the target point;
2. *starting realistic* – specifies a value close to the target point;
3. *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 3. 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 $v[is_l]_i^1$ of is_l :

$$starting_realistic(is_l, trg_{il}) = v[is_l]_i^1$$

where $v[is_l]_i^1 \in [trg_{il} - \epsilon, trg_{il} + \epsilon]$ and $\epsilon > 0$ is small. ■

Concession tactics are functions that compute new values for each issue is_l . The following five tactics are commonly discussed in the literature [9]:

1. *stalemate* – models a null concession on is_l ;
2. *tough* – models a small concession on is_l ;
3. *moderate* – models a moderate concession on is_l ;
4. *soft* – models a large concession on is_l ;
5. *accommodate* – models a complete concession on is_l .

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

Definition 4. 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 period 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 $Cf \in]0.10, 0.15]$).

3 Related Work

AI researchers have investigated the design of autonomous negotiating agents from two main perspectives: a theoretical or formal perspective and a practical or computational perspective. Researchers following the theoretical perspective attempt mainly to develop formal models for describing, specifying, and reasoning about the key features of the agents. Most researchers have focused on formal bargaining, auctions, market-oriented programming, contracting, and coalition formation (see, e.g., [5,8]). On the other hand, researchers following the practical perspective attempt mainly to develop

computational models for specifying the key data structures of the agents and the processes operating on these structures. Most models are based on informal procedures from the social sciences (see, e.g., [3,6,10]).

The task of designing and implementing negotiating agents involves the consideration of insights from multiple relevant research areas to integrate different perspectives on negotiation. In particular, game-theoretic (strategic) and behavioural negotiation theories can mutually reinforce each other and lead to more comprehensive and richer models. Yet, the majority of existing models largely ignore the integration of the results from both research areas. Noting this gap, this paper uses both game-theoretic techniques and methods from the social sciences as a basis to develop negotiating agents.

4 Conclusion

In this paper, we argue that an interdisciplinary approach towards the development of autonomous negotiating agents is possible and highly desirable. As a starting point for this research effort, we present a model that handles bilateral multi-issue negotiation, employs game-theoretic techniques to define equilibrium strategies for the bargaining game of alternating offers, and formalizes a set of negotiation strategies studied in the social sciences. Autonomous agents equipped with the model are currently being developed using the Jade framework. The agents are able to negotiate under both complete and incomplete information, thereby making the model in particular and the agents in general very compelling for automated negotiation.

Our aim for the future is to extend the model and to perform its experimental validation. In particular, the model defines a number of strategies based on rules-of-thumb distilled from behavioral negotiation theory. Hence, these strategies need to be empirically evaluated to determine precisely how they behave in different situations. In addition, we intend to study the bargaining game of alternating offers in order to define equilibrium strategies for two incompletely informed players.

References

1. Bichler, M., Kersten, G., Strecker, S.: Towards a Structured Design of Electronic Negotiations. *Group Decision and Negotiation* 12, 311–335 (2003)
2. Bussmann, S., Jennings, N., Wooldridge, M.: *Multiagent Systems for Manufacturing Control*. Springer, Heidelberg (2004)
3. Faratin, P.: *Automated Service Negotiation Between Autonomous Computational Agents*, Ph.D. Thesis, Queen Mary & Westfield College, UK (2000)
4. Fasli, M.: *Agent Technology for e-Commerce*. John Wiley & Sons, Chichester (2007)
5. Fatima, S., Wooldridge, M., Jennings, N.: Multi-Issue Negotiation with Deadlines. *Journal of Artificial Intelligence Research* 27, 381–417 (2006)
6. Ito, T., Hattori, H., Zhang, M., Matsuo, T.: *Rational, Robust, and Secure Negotiations in Multi-Agent Systems*. Springer, Heidelberg (2008)
7. Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Wooldridge, M., Sierra, C.: Automated Negotiation: Prospects, Methods and Challenges. *Group Dec. and Neg.* 10, 199–215 (2001)
8. Kraus, S.: *Strategic Negotiation in Multi-Agent Environments*. MIT Press, Cambridge (2001)

9. Lewicki, R., Barry, B., Saunders, D., Minton, J.: *Negotiation*. McGraw Hill, New York (2003)
10. Lopes, F., Mamede, N., Novais, A.Q., Coelho, H.: A Negotiation Model for Autonomous Computational Agents: Formal Description and Empirical Evaluation. *Journal of Intelligent & Fuzzy Systems* 12, 195–212 (2002)
11. Lopes, F., Mamede, N., Novais, A.Q., Coelho, H.: Negotiation Strategies for Autonomous Computational Agents. In: *ECAI 2004*, pp. 38–42. IOS Press, Amsterdam (2004)
12. Lopes, F., Novais, A.Q., Coelho, H.: Interdisciplinary Approach to Automated Negotiation: A Preliminary Report. In: *ECMS 2008*. SCS Publishing House (to appear, 2008)
13. Martello, S., Toth, P.: *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, Chichester (1990)
14. Osborne, M., Rubinstein, A.: *Bargaining and Markets*. Academic Press, London (1990)
15. Pruitt, D., Kim, S.: *Social Conflict: Escalation, Stalemate, Settlement*. McGraw Hill, New York (2004)
16. Rahwan, I., Ramchurn, S., Jennings, N., McBurney, P., Parsons, S., Sonenberg, L.: Argumentation-based Negotiation. *The Knowledge Engineering Review* 18, 343–375 (2004)
17. Raiffa, H.: *The Art and Science of Negotiation*. Harvard University Press (1982)
18. Simon, H.: *The Sciences of the Artificial*. MIT Press, Cambridge (1981)
19. Thompson, L.: *The Mind and Heart of the Negotiator*. Prentice-Hall, Englewood Cliffs (2005)