

Concession Behaviour in Automated Negotiation

Fernando Lopes¹ and Helder Coelho²

¹ LNEG – National Research Institute
Estrada do Paço do Lumiar 22, 1649-038 Lisbon, Portugal
fernando.lopes@ineti.pt

² University of Lisbon, Department of Computer Science
Bloco C6, Piso 3, Campo Grande, 1749-016 Lisbon, Portugal
hcoelho@di.fc.ul.pt

Abstract. Traditional negotiation, conducted face-to-face and via mail or telephone, is often difficult to manage, prone to misunderstanding, and time consuming. Automated negotiation promises a higher level of process efficiency, and more importantly, a faster emergence and a higher quality of agreements. The potential monetary impact has led to an increasing demand for systems composed of software agents representing individuals or organizations and capable of reaching efficient agreements. At present, work on automated negotiation has generated many useful ideas and concepts leading to important theories and systems. Yet, the design of software agents with negotiation competence largely lacks systematic, traceable, and reproducible approaches, and thus remains more an art than a science. Against this background, this paper presents a model for software agents that handles two-party and multi-issue negotiation. The model incorporates various concession strategies and negotiation tactics. Concession strategies are computationally tractable functions that define the tactics to be used both at the outset and throughout negotiation. Tactics, in turn, are functions that specify the short-term moves to be made at each point of negotiation.

Keywords: Automated negotiation, Negotiation strategies, Negotiation tactics, Multi-agent systems.

1 Introduction

Negotiation is a discussion among conflicting parties with the aim of reaching agreement about a divergence of interests [13]. The list of situations that can be handled by negotiation is endless. Some situations are purely competitive, as when the parties have completely opposed interests. Other situations are purely cooperative, as when the parties have perfectly compatible interests. Most situations are mixed-motive, containing elements of both competitive and cooperative situations – the parties' interests are imperfectly correlated [14]. There are, however, several characteristics common to most negotiation situations, including [6]: (i) two or more parties, (ii) a conflict among the parties, and (iii) an individual preference to search for agreement rather than to appeal to a higher authority, to permanently break off contact, or to fight openly.

Negotiation may involve two parties (bilateral negotiation) or more than two parties (multilateral negotiation) and one issue (single-issue negotiation) or many issues (multi-issue negotiation). Also, negotiation may proceed through several distinct phases or stages, notably a beginning or initiation phase, a middle or problem-solving phase, and an ending or resolution phase [6]. The initiation phase focuses on preparation and planning for negotiation – it is marked by each party’s efforts to emphasize points of difference and to posture for positions. The problem-solving phase seeks a solution for a dispute – it is characterized by extensive interpersonal interaction, strategic maneuvers, and movement toward a mutually acceptable agreement. The resolution phase focuses on details and implementation of a final agreement.

Traditional negotiation, conducted face-to-face and via mail or telephone, is often difficult to manage, prone to misunderstanding, and time consuming [1]. Negotiators are typically satisfied with the final outcome and, in many instances, proudly describe it. However, they frequently view conflict-laden situations with a fundamentally more distrustful, win-lose attitude than is necessary or desirable, and settle for outcomes that are worse for them than other available solutions [19]. They often fail to achieve agreements on the Pareto optimal or efficient frontier (*i.e.*, the locus of achievable joint evaluations from which no joint gains are possible [17]).

Automated negotiation promises a higher level of process efficiency, and most importantly, a faster emergence and a higher quality of agreements. The potential monetary impact has led to an increasing demand for systems composed of software agents representing individuals or organizations and capable of reaching mutually beneficial agreements (e.g., the industrial trend toward agent-based supply chain management). Yet, the design of software agents with negotiation competence largely lacks systematic, traceable, and reproducible approaches, and thus remains more an art than a science. There is much further work to be done, and some current ideas and concepts are likely to be substantially altered as researchers move ahead (but see [10]).

Against this background, this paper presents a model for software agents that handles two-party and multi-issue negotiation. The model incorporates a bilateral negotiation protocol, a set of concession strategies, and a set of negotiation tactics. The protocol formalizes the set of possible tasks that the agents can perform during the course of negotiation. The strategies and tactics formalize the tasks that each agent should perform to negotiate effectively. More specifically, the strategies define the tactics to be used both at the beginning and during the course of negotiation. The tactics formalize the individual moves to be made at each point of the negotiation process.

This paper builds on our previous work in the area of automated negotiation. In particular, it extends the work presented in [7,8,9] by introducing precise definitions for the key components of our model. Also, it formalizes concession strategies as computationally tractable functions that specify the tactics to be used both at the outset and throughout negotiation. Furthermore, at every period, strategies state whether bargaining should continue or terminate.

The remainder of the paper is structured as follows. Section 2 presents a negotiation model for software agents. Section 3 discusses related work and compares the negotiation model with other existing models. Finally, section 4 presents concluding remarks and indicates future avenues of research.

2 A Negotiation Model for Software Agents

Let $\mathcal{A} = \{a_1, a_2\}$ be the set of autonomous agents (negotiating parties). Both the number of agents and their identity are fixed and known to all the participants. The negotiation issues $\{x_1, \dots, x_n\}$ are quantitative in nature and defined over continuous domains $\{D_1, \dots, D_n\}$, respectively. For each issue x_k , the range of acceptable values is represented by the interval $D_k = [min_k, max_k]$. The issues are also known to all the participants.

Preparation and planning are often considered the foundations for success in negotiation [6]. Accordingly, effective negotiators often make efforts to perform a number of activities before starting to bargain, including: (i) prioritizing the issues, and (ii) defining realistic, pessimistic, and optimistic targets. Prioritization usually involves two steps, namely deciding which issues are most important and which are least important, and determining whether the issues are connected or separate. Priorities can be set in a number of ways (e.g., to use standard techniques, such as the nominal group technique). For the sake of simplicity, we consider that negotiators set priorities by ranking-order the issues.

Target setting usually involves defining three key points for each issue at stake:

1. *the resistance point or limit* – the point where negotiators decide to stop the negotiation rather than to continue, because any settlement beyond this point is not minimally acceptable;
2. *the target point or level of aspiration* – the point where negotiators realistically expect to achieve a settlement;
3. *the optimistic point or asking price* – the best deal negotiators could possibly hope to assume.

We present below precise definitions for these intuitions.

Definition 1 (Issue, Agenda). *A negotiation issue is a resource to be allocated or a consideration to be resolved in negotiation. The negotiating agenda is the set $\mathcal{I} = \{x_1, \dots, x_n\}$ of issues to be deliberated during negotiation.* ■

Definition 2 (Priority, Weight). *The priority prt_k^i of an agent $a_i \in \mathcal{A}$ for an issue $x_k \in \mathcal{I}$ is a number that represents the importance of x_k . The weight w_k^i is a number that represents the preference for x_k .* ■

Definition 3 (Limit, Target Point, Optimistic Point). *The limit lim_k^i of an agent $a_i \in \mathcal{A}$ for an issue $x_k \in \mathcal{I}$ is the ultimate fallback position for x_k , the point beyond which a_i is unwilling to concede on x_k . The target point trg_k^i is the point at which a_i is satisfied with the value of x_k . The optimistic point opt_k^i is the most preferred or ideal value for x_k .* ■

2.1 The Negotiation Protocol and Negotiators' Preferences

The negotiation protocol is an alternating offers protocol [11]. Two agents or players bargain over the division of the surplus of $n \geq 2$ distinct issues. The players determine an allocation of the issues by alternately submitting proposals at times in $\mathcal{T} = \{1, 2, \dots\}$. This means that one proposal is made per time period $t \in \mathcal{T}$, with an agent, say $a_i \in \mathcal{A}$, offering in odd periods $\{1, 3, \dots\}$, and the other agent $a_j \in \mathcal{A}$ offering in even periods $\{2, 4, \dots\}$. The agents have the ability to unilaterally opt out of the negotiation when responding to a proposal.

The negotiation process starts with a_i submitting a proposal $p_{i \rightarrow j}^1$ to a_j in period $t=1$. The agent a_j receives $p_{i \rightarrow j}^1$ and can either accept the offer (**Yes**), reject it and opt out of the negotiation (**Opt**), or reject it and continue bargaining (**No**). In the first two cases the negotiation ends. Specifically, if $p_{i \rightarrow j}^1$ is accepted, negotiation ends successfully and the agreement is implemented. Conversely, if $p_{i \rightarrow j}^1$ is rejected and a_j decides to opt out, negotiation terminates with no agreement. In the last case, negotiation proceeds to the next time period $t=2$, in which a_j makes a counter-proposal $p_{j \rightarrow i}^2$. The tasks just described are then repeated. Once an agreement is reached, the agreed-upon allocations of the issues are implemented.

Definition 4 (Proposal). Let \mathcal{A} be the set of negotiating agents and \mathcal{I} the set of issues at stake in negotiation. Let \mathcal{T} be the set of time periods. A proposal $p_{i \rightarrow j}^t$ submitted by an agent $a_i \in \mathcal{A}$ to an agent $a_j \in \mathcal{A}$ in period $t \in \mathcal{T}$ is a vector of issue values:

$$p_{i \rightarrow j}^t = (v_1, \dots, v_n)$$

where v_k , $k = 1, \dots, n$, is a value of an issue $x_k \in \mathcal{I}$. ■

Definition 5 (Agreement, Possible Agreements). An agreement is a proposal accepted by all the negotiating agents in \mathcal{A} . The set of possible agreements is:

$$\mathcal{S} = \{(v_1, \dots, v_n) \in \mathbb{R}^n : v_k \in D_k, \text{ for } k = 1, \dots, n\}$$

where v_k is a value of an issue $x_k \in \mathcal{I}$. ■

Negotiators should express their own preferences to rate and compare incoming offers and counter-offers. The most common way to model the preferences of the negotiating agents is probably to define a utility function over all possible outcomes [4,17]. Let $\mathcal{I} = \{x_1, \dots, x_n\}$ be the agenda and $\mathcal{D} = \{D_1, \dots, D_n\}$ the set of issue domains. We consider that each agent $a_i \in \mathcal{A}$ has a continuous utility function: $U_i : \{D_1 \times \dots \times D_n\} \cup \{\text{Opt}, \text{Disagreement}\} \rightarrow \mathbb{R}$. Accordingly, when the utility for a_i from one outcome is greater than from another outcome, we assume that a_i prefers the first outcome over the second. The outcome **Opt** is interpreted as one of the agents opting out of the negotiation in a given period of time. Perpetual disagreement is denoted by **Disagreement**.

Now, the additive model is probably the most widely used in multi-issue negotiation – the parties assign numerical values to the different levels on each issue and add them to get an entire offer evaluation [17]. This model is simple

and intuitive, and therefore well suited to the purposes of this work. We consider that each agent a_i has a scoring or single-issue (marginal) utility function for each issue at stake in negotiation, *i.e.*, a function that gives the score a_i assigns to a value of an issue x_k . For convenience, scores are kept in the interval $[0,1]$. Additionally, as mentioned above, we consider that a_i has a multi-issue utility function to rate offers.

Definition 6 (Multi-Issue Utility Function). *Let $\mathcal{A} = \{a_1, a_2\}$ be the set of negotiating agents and $\mathcal{I} = \{x_1, \dots, x_n\}$ the negotiating agenda. The utility function U_i of an agent $a_i \in \mathcal{A}$ to rate offers and counter-offers takes the form:*

$$U_i(x_1, \dots, x_n) = \sum_{k=1}^n w_k^i \times V_k^i(x_k)$$

where:

- (i) w_k^i is the weight of a_i for an issue $x_k \in \mathcal{I}$;
- (ii) $V_k^i(x_k)$ is the (marginal) utility function of a_i for x_k .

■

Negotiation may end with either agreement or no agreement. Failure to agree can occur in two ways: (i) either party decides to opt out unilaterally, or (ii) the two do not agree to any proposal. The resistance points or limits play a key role in reaching agreement when the parties have the ability to unilaterally opt out of the negotiation – they define the worst agreement for a given party which is still better than opting out. For each agent $a_i \in \mathcal{A}$, we will denote this agreement by $\hat{s}_i \in \mathcal{S}$. Hence, \hat{s}_i will be the least-acceptable agreement for a_i , *i.e.*, the worst (but still acceptable) agreement for a_i . The set of all agreements that are preferred by a_i to opting out will be denoted by S_i .

Definition 7 (Least-acceptable Agreement, Acceptable Agreements). *The least-acceptable agreement for an agent $a_i \in \mathcal{A}$ is defined as: $\hat{s}_i = (\lim_1^i, \dots, \lim_n^i)$, where \lim_k^i , $k = 1, \dots, n$, is the limit of a_i for an issue $x_k \in \mathcal{I}$. The set of acceptable agreements for a_i is:*

$$S_i = \{s: s \in \mathcal{S}, U_i(s) \geq U_i(\hat{s}_i)\}$$

where $U_i(\hat{s}_i)$ is the utility of \hat{s}_i for a_i .

■

Perpetual disagreement is the least-preferred or worst outcome, *i.e.*, disagreement is even worse than opting out. Thus, the agents prefer any agreement in any given time period over the continuation of the negotiation process indefinitely. Formally, and more precisely, we state the following:

- (1) (Acceptable agreements versus opting out). For every agent $a_i \in \mathcal{A}$ and acceptable agreement $s \in S_i$, $U_i(s) \geq U_i(\text{Opt})$.
- (2) (Opting out versus Disagreement). For every agent $a_i \in \mathcal{A}$, $U_i(\text{Opt}) > U_i(\text{Disagreement})$.

2.2 Concession Strategies

The opening offer and the initial concessions are two central elements of negotiation [19]. When negotiation begins, the parties are faced with a fundamental question – should the opening offer be exaggerated, more toward the optimistic point, or modest, somewhat closer to the limit? The main advantages of an exaggerated initial offer are [12]: (i) negotiators can concede further and hence elicit more counterconcessions from their opponent, and (ii) negotiators' later demands are likely to look generous. However, an exaggerated opening offer frequently communicates an attitude of toughness that may be harmful to long-term relationships. Also, it may be seen as too high by the other party and therefore summarily rejected. By contrast, an opening offer seen as reasonable or modest by the other party could perhaps have been higher, either to leave more room for movement or to achieve a higher settlement.

After the first round of offers, another fundamental question is, what concessions are to be made? Negotiators can choose to make none, holding firm and insisting on their original positions. By taking a firm position, negotiators attempt to capture most of the initial bargaining or settlement range (defined by the opening offers of both parties). However, there is the very real possibility that firmness will be reciprocated – one or both parties may become intransigent and withdraw completely. Negotiators can also choose to make some concessions, being flexible and changing their original positions. Flexibility often keeps negotiation going – the more flexible one party seems to be, the more the other party will believe that a settlement is possible. Hence, if concessions are to be made, another fundamental question is, how large should they be?

Concession strategies are computationally tractable functions that model significant opening positions and typical patterns of concessions. They specify the tactics to be used at the outset of negotiation (to prepare the initial offers). Also, at each step of negotiation, they specify the tactics to be used in preparing counter-offers. Furthermore, concession strategies state whether bargaining should continue or terminate. The words “computationally tractable functions” presume that agents are able to compute concession strategies in a reasonable amount of time. A formal definition of a generic strategy follows.

Definition 8 (Concession Strategy). Let \mathcal{A} be the set of negotiating agents, \mathcal{I} the negotiating agenda, \mathcal{T} the set of time periods, and \mathcal{S} the set of possible agreements. Let $a_i \in \mathcal{A}$ be the first agent to submit a proposal and T_i his set of tactics. A concession strategy $C_i : T_i \times T_i \times \mathcal{T} \rightarrow \mathcal{S} \cup \{\text{Yes}, \text{No}, \text{Opt}\}$ for a_i is a function with the following general form:

$$C_i = \begin{cases} \text{apply } O_i(x_k) \text{ and offer } p_{i \rightarrow j}^1, & \text{if } a_i \text{'s turn and } t=1 \\ \text{reject } p_{j \rightarrow i}^{t-1} \text{ and quit,} & \text{if } a_j \text{'s turn and } U_i(p_{j \rightarrow i}^{t-1}) < U_i(\hat{s}_i) \\ \text{apply } Y_i(x_k, f_k^i) \text{ and prepare } p_{i \rightarrow j}^t & \text{if } a_j \text{'s turn and } U_i(p_{j \rightarrow i}^{t-1}) \geq U_i(\hat{s}_i) \\ \text{if } U_i^* \geq 0 \text{ accept } p_{j \rightarrow i}^{t-1} \text{ else reject,} & \\ \text{offer compromise } p_{i \rightarrow j}^t, & \text{if } a_i \text{'s turn and } t > 1 \end{cases}$$

where:

- (i) $p_{i \rightarrow j}^1$ is the opening offer of a_i , $p_{j \rightarrow i}^{t-1}$ is the offer of a_j for period $t-1$ of negotiation, and $p_{i \rightarrow j}^t$ is the offer of a_i for the next period t of negotiation;
- (ii) for each issue $x_k \in \mathcal{I}$, $O_i(x_k)$ is an opening negotiation tactic, $Y_i(x_k, f_k^i)$ is a concession tactic, and $f_k^i \in [0, 1]$ is a real number that defines the magnitude of a concession on x_k , referred to as the concession factor of a_i for x_k (see subsection 2.3, below);
- (iii) $U_i(\hat{s}_i)$ is the utility of the least-acceptable agreement for a_i ;
- (iv) $U_i^* = U_i(p_{j \rightarrow i}^{t-1}) - U_i(p_{i \rightarrow j}^t)$

Two explanatory and cautionary notes are in order here. First, notation is being abused somewhat, by using C_i rather than $C_i(O_i(x_k), Y_i(x_k, f_k^i), t)$. The abuse helps improve readability, however, and meaning will always be clear from context. Second, tactics are functions of a single issue rather than a vector of issues. This permits great flexibility, since it allows agents to model a wide range of concession behaviors (e.g., large initial demands and slow concession making).

Interestingly, bargainers sometimes have different strengths of preference for the issues at stake – they place greater emphasis on some key issues and make significant efforts to resolve them favourably. Hence, they often yield on less important or low-priority issues, in the hope that their opponent will make compensating concessions [17,19]. A generic low-priority concession making strategy for a_i takes the form (again, the definition slightly abuses notation):

$$C_{LP}^i = \begin{cases} \text{apply } O_i \text{ and offer } p_{i \rightarrow j}^1, & \text{if } a_i \text{'s turn and } t=1 \\ \text{reject } p_{j \rightarrow i}^{t-1} \text{ and quit,} & \text{if } a_j \text{'s turn and } U_i(p_{j \rightarrow i}^{t-1}) < U_i(\hat{s}_i) \\ \text{apply } Y_i \text{ to } \hat{I}_i \text{ and prepare } p_{i \rightarrow j}^t & \text{if } a_j \text{'s turn and } U_i(p_{j \rightarrow i}^{t-1}) \geq U_i(\hat{s}_i) \\ \text{if } U_i^* \geq 0 \text{ accept } p_{j \rightarrow i}^{t-1} \text{ else reject,} & \\ \text{offer compromise } p_{i \rightarrow j}^t, & \text{if } a_i \text{'s turn and } t > 1 \end{cases}$$

where $\hat{I}_i \subset \mathcal{I}$ is the set of issues that are of lower priority to a_i . The definition of a specific strategy involves basically the specification of a particular opening negotiation tactic (e.g., the tactic “**starting optimistic**”) and a key concession tactic to apply to the issues of low priority (e.g., the tactic “**moderate**”). These and other negotiation tactics are defined in the next subsection.

2.3 Negotiation Tactics

Negotiation tactics are functions that model the short-term moves designed to enact high-level strategies. The following two groups of tactics will receive the preponderance of our attention in this paper:

1. *opening negotiation tactics* – functions that specify the demands to be made at the outset of negotiation;
2. *concession tactics* – functions that model the concessions to be made throughout negotiation.

As mentioned above, tactics are functions of a single issue.

Opening Negotiation Tactics. Skilled negotiators often start with high demands to leave room for later movement and hence elicit counterconcessions from their opponent [12]. High initial demands also protect limits from detection and underestimation (this is a concern about image loss). If limits are detected by the opponent, he may become unwilling to accept a better offer than the least-acceptable one, dooming all higher aspirations. If limits are underestimated, the opponent may become committed to unacceptable demands, fostering breakdown of negotiation. Thus, to avoid these dual dangers, bargainers typically place their demands well above their limits as a sort of smoke screen. Furthermore, high initial demands are also partly designed to protect target points (this is a concern about position loss). Clearly, bargainers often need to move in concert with their opponent toward mutually acceptable agreements. This means starting higher than targets and only moving down to them in coordination with the opponent [14].

Noticeably, starting high frequently communicates an attitude of toughness that can be reciprocated by the opponent, thus making negotiation “difficult to resolve” [6]. Hence, should bargainers start with a firm, determined stance, or adopt a position of moderateness and understanding? It follows that bargainers often decide how much to demand on the basis of the concessions they expect from their opponent – the farther the opponent is expected to concede, the more will be demanded (this phenomenon is referred to as tracking).

In general, three levels of initial demand are commonly discussed in the negotiation literature [6,15]: extreme or high, reasonable or moderate, and modest or low. They have motivated the definition of the following opening negotiation tactics:

1. *starting optimistic* – specifies a value for an issue close to the optimistic point;
2. *starting realistic* – specifies a value for an issue in the range defined by the target and the optimistic points;
3. *starting pessimistic* – specifies a value for an issue in the range defined by the target and the resistance points.

A formal definition of the tactic “*starting optimistic*” follows.

Definition 9 (Starting Optimistic). Let $\mathcal{A} = \{a_1, a_2\}$ be the set of negotiating agents and $\mathcal{I} = \{x_1, \dots, x_n\}$ the negotiating agenda. Let $\mathcal{D} = \{D_1, \dots, D_n\}$ be the set of issue domains. The tactic *starting optimistic* of an agent $a_i \in \mathcal{A}$ for an issue $x_k \in \mathcal{I}$ takes the form:

$$O_i(x_k) = opt_k^i + \epsilon$$

where:

- (i) $\epsilon > 0$ is small;
- (ii) opt_k^i is the optimistic point of a_i for x_k . ■

The definition of the other two tactics is essentially identical to that of “*starting optimistic*”, and is therefore omitted.

Concession Tactics. Concessions are a powerful aspect of negotiation – without them, in fact, some researchers consider that negotiation would not exist [19]. A concession is usually defined as a change of offer in the supposed direction of the other party’s interests that reduces the level of benefit sought. Concession rate is the speed at which demand level declines over time. A bargainer’s demand level can be thought of as the level of benefit to the self associated with the current demand or offer [12,15].

Practically speaking, bargainers often enter negotiation expecting concessions. Their opening position may be good for both sides and might have been the final settlement if the parties started negotiation from different points. Even so, bargainers generally resent a take-it-or-leave-it approach – an offer that may have been accepted had it emerged as a result of concession making may be rejected when it is thrown on the table and presented as a *fait accompli* [6]. Ample research evidence indicates that the parties feel better about a settlement when negotiation has involved a progression of concessions [14].

A formal definition of a generic concession tactic follows (in the interests of readability, and without loss of generality, we consider that $a_i \in \mathcal{A}$ wants to maximize $x_k \in \mathcal{I}$).

Definition 10 (Concession Tactic). Let $\mathcal{A} = \{a_1, a_2\}$ be the set of negotiating agents, $\mathcal{I} = \{x_1, \dots, x_n\}$ the negotiating agenda, and $\mathcal{D} = \{D_1, \dots, D_n\}$ the set of issue domains. A concession tactic $Y_i: D_k \times [0, 1] \rightarrow D_k$ of an agent $a_i \in \mathcal{A}$ for an issue $x_k \in \mathcal{I}$ is a function with the following general form:

$$Y_i(x_k, f_k^i) = x_k - f_k^i(x_k - \lim_k^i)$$

where:

- (i) f_k^i is the concession factor of a_i for x_k ;
- (ii) \lim_k^i is the limit of a_i for x_k . ■

Negotiators may consider strikingly different patterns of concessions as negotiation unfolds. However, the following three levels of concession magnitude are commonly discussed in the negotiation literature [6,14]: large, substantial, and small. To this we would add two other levels: null and complete. Accordingly, we consider the following five concession tactics:

1. *stalemate* – models a null concession on an issue x_k at stake;
2. *tough* – models a small concession on x_k ;
3. *moderate* – models a substantial concession on x_k ;
4. *soft* – models a large concession on x_k ;
5. *accommodate* – models a complete concession on x_k .

These and other similar tactics are defined by considering specific values for the concession factor f_k^i . In particular, the “*stalemate*” tactic is defined by $f_k^i = 0$ and the “*accommodate*” tactic by $f_k^i = 1$. The other three tactics are defined by considering values for f_k^i in different ranges (e.g., the “*tough*” tactic by $f_k^i \in]0.00, 0.05]$, the “*moderate*” tactic by $f_k^i \in]0.05, 0.15]$, and the “*soft*” tactic by $f_k^i \in]0.15, 0.20]$).

3 Related Work

Artificial intelligence (AI) researchers have investigated the design of agents with negotiation competence from two main perspectives: a theoretical or formal mathematical perspective and a practical or system-building perspective. Researchers following the theoretical perspective have attempted mainly to develop formal models of negotiation, *i.e.*, models for describing, specifying, and reasoning about the key features of negotiating agents. To this end, they have drawn heavily on game-theoretic and economic methods (see, e.g., [2,5,18]). On the other hand, researchers following the practical perspective have attempted mainly to develop computational models of negotiation, *i.e.*, models for specifying the key data structures of negotiating agents and the processes operating on these structures. They have drawn heavily on social sciences techniques for understanding interaction and negotiation (see, e.g., [3,10,16]).

Overall, various researchers have developed models that incorporate specific protocols (notably, the alternating offers protocol) and libraries of negotiation strategies (notably, concession strategies). However, the authors are aware of no similar efforts to define strategies as functions that specify the tactics to be used both at the outset and throughout negotiation. Tactics, in turn, are defined as functions that specify the short-term moves to be made at each point of negotiation. Our interest lies mainly in formalizing important strategies and tactics motivated by rules-of-thumb distilled from good behavioral practice in real-life negotiations.

4 Conclusion

This paper has presented a model for software agents that handles two-party and multi-issue negotiation. The model incorporates a bilateral negotiation protocol, a set of concession strategies, and a set of negotiation tactics. The protocol is an alternating offers protocol. The strategies are computationally tractable functions that define the tactics to be used both at the beginning and during the course of negotiation. The words “computationally tractable functions” presume that agents are able to compute the strategies in a reasonable amount of time. Furthermore, at every period of negotiation, the strategies state whether bargaining should continue or terminate. The tactics are functions that specify the individual moves to be made at each point of the negotiation process.

Autonomous agents equipped with the negotiation model are currently being developed. Our aim for the future is to perform a number of experiments to empirically evaluate the key components of the agents. Also, notice that the task of designing and implementing agents with negotiation competence involves the consideration of insights from multiple relevant research areas. Accordingly, we also intend to develop an interdisciplinary framework for automated negotiation – game-theoretic (strategic) and behavioural negotiation theories should mutually reinforce each other and contribute to richer negotiators.

References

1. Bichler, M., Kersten, G., Strecker, S.: Towards a Structured Design of Electronic Negotiations. *Group Decision and Negotiation* 12, 311–335 (2003)
2. Fatima, S., Wooldridge, M., Jennings, N.: A Comparative Study of Game Theoretic and Evolutionary Models of Bargaining for Software Agents. *Artificial Intelligence Review* 23, 185–203 (2005)
3. Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Wooldridge, M., Sierra, C.: Automated Negotiation: Prospects, Methods and Challenges. *Group Decision and Negotiation* 10, 199–215 (2001)
4. Keeney, R., Raiffa, H.: Decisions with Multiple Objectives: Preferences and Value Tradeoffs. John Wiley & Sons, Chichester (1976)
5. Kraus, S.: Strategic Negotiation in Multi-Agent Environments. MIT Press, Cambridge (2001)
6. Lewicki, R., Barry, B., Saunders, D., Minton, J.: Negotiation. McGraw Hill, New York (2003)
7. 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)
8. Lopes, F., Mamede, N., Novais, A.Q., Coelho, H.: Negotiation Strategies for Autonomous Computational Agents. In: 16th European Conference on Artificial Intelligence (ECAI-04), pp. 38–42. IOS Press, Amsterdam (2004)
9. Lopes, F., Mamede, N., Novais, A.Q., Coelho, H.: Negotiation Among Autonomous Agents: Experimental Evaluation of Integrative Strategies. In: 12th Portuguese Conference on Artificial Intelligence, pp. 280–288. IEEE Computer Society Press, Los Alamitos (2005)
10. Lopes, F., Wooldridge, M., Novais, A.Q.: Negotiation Among Autonomous Computational Agents: Principles, Analysis and Challenges. *Artificial Intelligence Review* 29, 1–44 (2008)
11. Osborne, M., Rubinstein, A.: Bargaining and Markets. Academic Press, San Diego (1990)
12. Pruitt, D.: Negotiation Behavior. Academic Press, New York (1981)
13. Pruitt, D.: Social Conflict. In: Gilbert, D., Fiske, S., Lindzei, G. (eds.) *The Handbook of Social Psychology*, vol. 2, pp. 470–503. McGraw-Hill, New York (1998)
14. Pruitt, D., Carnevale, P.: Negotiation in Social Conflict. Open University Press, Philadelphia (1993)
15. Pruitt, D., Kim, S.: Social Conflict: Escalation, Stalemate, and 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, Cambridge (1982)
18. Sandholm, T.: Distributed Rational Decision Making. In: Weiss, G. (ed.) *Multi-Agent Systems – A Modern Approach to Distributed Artificial Intelligence*, pp. 201–259. MIT Press, Cambridge (1999)
19. Thompson, L.: The Mind and Heart of the Negotiator. Prentice-Hall, Englewood Cliffs (2005)