

## Negotiation decision functions for autonomous agents

Peyman Faratin<sup>\*</sup>, Carles Sierra<sup>1</sup>, Nick R. Jennings<sup>2</sup>

*Department of Electronic Engineering, Queen Mary and Westfield College, University of London, London E1 4NS, UK*

Received 22 October 1997; received in revised form 16 February 1998; accepted 23 March 1998

---

### Abstract

We present a formal model of negotiation between autonomous agents. The purpose of the negotiation is to reach an agreement about the provision of a service by one agent for another. The model defines a range of strategies and tactics that agents can employ to generate initial offers, evaluate proposals and offer counter proposals. The model is based on computationally tractable assumptions, demonstrated in the domain of business process management and empirically evaluated. © 1998 Elsevier Science B.V. All rights reserved.

*Keywords:* Multi-agent systems; Automated negotiation; Business process management

---

### 1. Introduction

Autonomous agents are being increasingly used in a wide range of industrial and commercial domains [2]. These agents have a high degree of self-determination – they decide for themselves what, when and under what conditions their actions should be performed. In most cases, such agents need to interact with other autonomous agents to achieve their objectives (either because they do not have sufficient capabilities or resources to complete their problem solving alone or because there are interdependencies between the agents). The objectives of these interactions are to make other agents undertake a particular course of action (e.g. perform a particular service), modify a planned course of action (e.g. delay or bring forward a particular action so that there is no longer a conflict), or come to an agreement on a common course of action. Since the agents have no direct control over one another, they must persuade their acquaintances to act in particular ways (they cannot simply instruct them). The type of persuasion we consider is negotiation – *a process by which a joint decision is made by two or more parties. The parties first verbalise contradictory demands and then move towards agreement by a process of concession making or search for new alternatives* [11].

Given its pervasive nature, negotiation comes in many shapes and forms. However in this work we are interested in a particular class of negotiation – namely *service-oriented negotiation*. In this context, one agent (the client) requires

---

<sup>\*</sup> Corresponding author. E-mail: p.faratin@qmw.ac.uk

<sup>1</sup> On sabbatical leave from Artificial Intelligence Research Institute – IIIA, Spanish Council for Scientific Research – CSIC. 08193 Bellaterra, Barcelona, Spain. E-mail: sierra@iia.csic.es. With the support of the Spanish Ministry of Education grant PR95-313.

<sup>2</sup> E-mail: n.r.jennings@qmw.ac.uk

a service to be performed on its behalf by some other agent (the server).<sup>3</sup> Negotiation involves determining a contract under certain terms and conditions. The negotiation may be iterative in that several rounds of offers and counter offers will occur before an agreement is reached or the negotiation is terminated.

When building an autonomous agent which is capable of flexible and sophisticated negotiation, three broad areas need to be considered [10] – what negotiation protocol will be used?, what are the issues over which negotiation takes place?, and what reasoning model will the agents employ? This paper concentrates predominantly on the final point, although the protocol and negotiation object are briefly defined. A comprehensive reasoning model for service-oriented negotiation should determine: which potential agent should be contacted, whether negotiation should proceed in parallel with all agents or whether it should run sequentially, what initial offers should be sent out, what the range of acceptable agreements is, what counter offers should be generated, when negotiation should be abandoned, and when an agreement is reached.

To this end, this paper presents a formal account of a negotiating agent's reasoning component – in particular it concentrates on the processes of evaluating incoming proposals and generating counter proposals. The model specifies the key structures and processes involved in this endeavour and defines their inter-relationships. The model was shaped by practical considerations and insights emanating from the development of a system of negotiating agents for business process management (see [6] and Section 2 for more details). The main contributions of our model are: (i) it allows rich and flexible negotiation schemes to be defined; (ii) it is based on assumptions which are realistic for autonomous, computational agents (see Section 3.2 for the set of requirements and Section 7 for a discussion of related approaches) and (iii) its main properties have been empirically evaluated (see Section 6). In addition, we have some initial results on the convergence of certain types of negotiation using our model (although this aspect is not discussed in this paper, refer to [15] for more details).

In this paper we concentrate on many-parties, many-issues, single-encounter negotiations with an environment of limited resources (time among them). Section 2 gives details of the types of applications and scenarios in which we are interested. Sections 3–5 present the proposed model and in Section 6 the model is empirically evaluated. Finally, related work and some future avenues of research are outlined in Sections 7 and 8, respectively.

## 2. Service-oriented negotiation

This section characterises the context in which our service-oriented negotiations take place. The scenario is motivated by work in the ADEPT project [6] which developed generic negotiating agents for business process management applications. However, to provide a context for this work, one of the ADEPT applications, for managing a British Telecom (BT) process, is presented in detail (Section 2.1). This scenario is then analysed in terms of its key characteristics and assumptions as they relate to the process of negotiation (Section 2.2).

### 2.1. BT's provide customer quote business process

This scenario is based on BT's business process of providing a quotation for designing a network to provide particular services to a customer (Fig. 1).<sup>4</sup> The overall process receives a customer service request as its input and generates as its output a quote specifying how much it would cost to build a network to realise that service. It involves up to six agent types: the sales department agent, the customer service division agent, the legal department agent, the design division agent, the surveyor department agent, and the various agents who provide the out-sourced service of

<sup>3</sup> A service is a problem solving activity which has clearly defined start and end points. Examples include diagnosing a fault, buying a group of shares in the stock market, or allocating bandwidth to transmit a video-conference.

<sup>4</sup> The negotiations between the agents are denoted by arrows (arrow head toward client) and the service involved in the negotiation is juxtaposed to the respective arrow.

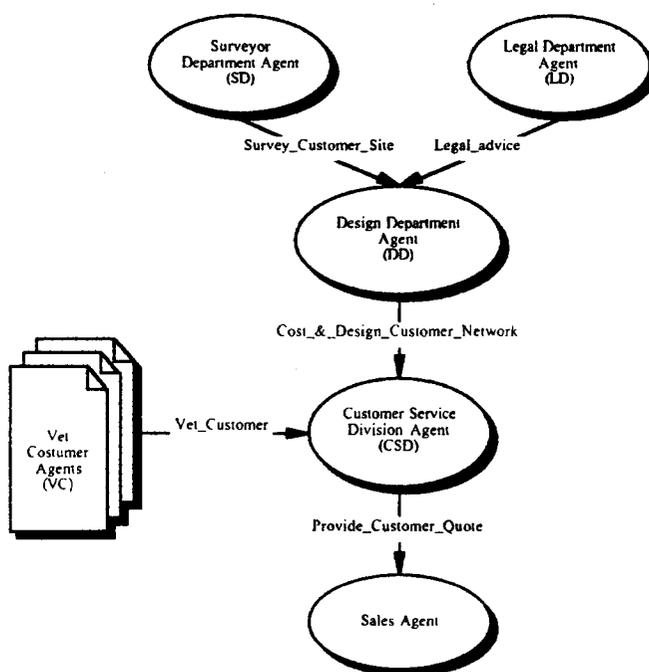


Fig. 1. Agent system for BT's provide customer quote business process.

vetting customers. All negotiations are centred on a multi-attribute object, where attributes are, for instance, price, quality, duration etc. of a service (see [6]).

The process is initiated by the sales agent which negotiates with the CSD agent (mainly over time, but also over the number of invocations and the form in which the final result should be delivered) for the service of providing a customer quote. The first stages of the *Provide-Customer-Quote* service involve the CSD agent capturing the customer's details and vetting the customer in terms of their credit worthiness. The latter sub-service is actually performed by one of the VC agents. Negotiation is used to determine which VC agent should be selected – the main attributes negotiated over are the price of the service, the penalty for contract violation, the desired quality of the service and the time by which the service should be performed. If the customer fails the vetting procedure, then the quote process terminates. Assuming the customer is satisfactory, the CSD agent maps their requirements against a service portfolio. If the requirements can be met by a standard off-the-shelf portfolio item, then an immediate quote can be offered based on previous examples. In the case of bespoke services, however, the process is more complex. The CSD agent negotiates with the DD agent (over time and quality) for the service of designing and costing the desired network service. In order for the DD agent to provide this service, it must negotiate with the LD agent (over time) and perhaps with the SD agent. The LD agent checks the design to ensure the legality of the proposed service (e.g. it is illegal to send unauthorised encrypted messages across France). If the desired service is illegal, then the entire quote process terminates and the customer is informed. If the requested service is legal, then the design phase can start. To prepare a network design, it is usually necessary to have a detailed plan of the existing equipment at the customer's premises. Sometimes such plans might not exist and sometimes they may be out of date. In either case, the DD agent determines whether the customer site(s) should be surveyed. If such a survey is warranted, the DD agent negotiates with the SD agent (over price and time) for the *Survey-Customer-Site* service. On completion of the network design and costing, the DD agent informs the CSD agent, which informs the customer of the service quote. The business process then terminates.

## 2.2. Characteristics and assumptions

The following negotiation characteristics can be observed in the aforementioned scenario. Moreover, it is believed that these characteristics are likely to be common to a wide range of service-oriented negotiations between autonomous agents.

- A given service can be provided by more than one agent (e.g. multiple agents can provide the vet customer service to the CSD agent). The available services may be identical in their characteristics or they may vary along several dimensions (e.g. quality, price, availability, etc.).
- Individual agents can be both clients and servers for different services in different negotiation contexts.
- Negotiations can range over a number of issues (e.g. price, duration, cost, etc.). Each successful negotiation requires a range of such issues to be resolved to the satisfaction of both parties. Agents may be required to make trade-offs between issues (e.g. faster completion time for lower quality) in order to come to an agreement.
- The social context and inter-relationships of the participants influences the way agents negotiate. Some negotiations involve entities within the same organisation (e.g. between the CSD and DD agents) and hence are generally cooperative in nature. Other negotiations are inter-organisational and purely competitive – involving self-interested, utility maximising agents (e.g. between the VC agents and the CSD agent). Some groups of agents often negotiate with one another for the same service (e.g. the CSD and DD agents), whereas other negotiations are more open in nature (for example, the set of VC agents changes frequently and hence the CSD agent often negotiates with unknown agents).
- As the agents are autonomous, the factors which influence their negotiation stance and behaviour are private and not available to their opponents (especially in inter-organisational settings). Thus agents do not know what utilities their opponents place on various outcomes, they do not know what reasoning models they employ, they do not know their opponent's constraints and they do not know whether an agreement is even possible at the outset (i.e. the participants may have non-intersecting ranges of acceptability).
- Since negotiation takes place within a highly intertwined web of activity (the business process), time is a critical factor. Timings are important on two distinct levels: (i) the time it takes to reach an agreement must be reasonable; and (ii) the time by which the negotiated service must be executed is important in most cases and crucial in others. The former means that the agents should not become involved in unnecessarily complex and time consuming negotiations – the time spent negotiating should be reasonable with respect to the value of the service agreement. The latter means that the agents sometimes have hard deadlines by which agreements must be in place (this occurs mainly when multiple services need to be combined or closely coordinated).

## 3. The negotiation model

The negotiation model in this section is based on a variation of the *two parties, many issues* value scoring system presented in [12]. That is, a model for bilateral negotiations about a set of quantitative variables. Our variation transforms that model into a *many parties, many issues* model (that is, multilateral negotiations about a set of variables). This is important since multilateral negotiations are common in the application domains in which we are interested. Our model of multilateral negotiations is based on a set of mutually influencing *two parties, many issues* negotiations. We call the sequence of offers and counter-offers in a two-party negotiation a *negotiation thread*. Offers and counter offers are generated by linear combinations of simple functions, called *tactics*. Tactics generate an offer, or counter offer, for a single component of the negotiation object (or issue) using a single criterion (time, resources, etc.). Different weights in the linear combination allow the varying importance of the criteria to be modelled. For example, when determining values of an issue, it may initially be more important to take into account the other agent's behaviour than the remaining time. In which case, the tactics that emphasize the behaviour of other agents will be given greater precedence than the tactics which base their value on the amount of time remaining.

However, to achieve flexibility in negotiation, the agents may wish to change their ratings of the importance of the different criteria over time. For example, remaining time may become correspondingly more important than the imitation of the other's behaviour as the time by which an agreement must be in place approaches. We use the term *strategy* to denote the way in which an agent changes the weights of the different tactics over time. Thus, strategies combine tactics depending on the history of negotiations and the internal reasoning model of the agents, and negotiation threads influence one another by means of strategies (see Section 5).

Before presenting our model, we introduce a basic, multi-attribute model for bilateral negotiation [7].

### 3.1. A bilateral negotiation model

Let  $i$  ( $i \in \{a, b\}$ ) represent the negotiating agents and  $j$  ( $j \in \{1, \dots, n\}$ ) the issues under negotiation. The set of issues in real world negotiations is always finite. Let  $x_j \in [\min_j^i, \max_j^i]$  be a value for issue  $j$  acceptable by agent  $i$ . Here we limit ourselves to considering issues for which negotiation amounts to determining a value between an agent's defined delimited range. Each agent has a scoring function  $V_j^i : [\min_j^i, \max_j^i] \rightarrow [0, 1]$  that gives the score agent  $i$  assigns to a value of issue  $j$  in the range of its acceptable values. For convenience, scores are kept in the interval  $[0, 1]$ .

The next element of the model is the relative importance that an agent assigns to each issue under negotiation.  $w_j^i$  is the importance of issue  $j$  for agent  $i$ . We assume that the weights of both agents are normalised, i.e.  $\sum_{1 \leq j \leq n} w_j^i = 1$ , for all  $i$  in  $\{a, b\}$ . With these elements in place, it is now possible to define an agent's scoring function<sup>5</sup> for a *contract* – that is, for a value  $x = (x_1, \dots, x_n)$  in the multi-dimensional space defined by the issues' value ranges:

$$V^i(x) = \sum_{1 \leq j \leq n} w_j^i V_j^i(x_j).$$

If both negotiators use such an additive scoring function, Raiffa showed that it is possible to compute the optimum value of  $x$  as an element on the efficient frontier of negotiation<sup>6</sup> (see [12, p. 164]).

For example, the set of negotiation issues for a server agent  $a$  may consist of {price, volume} – the price required to provide the service and the number of service instances attainable by  $a$ . In addition to this, let  $a$  have the following reservation values  $[\min_{\text{price}}^a, \max_{\text{price}}^a] = [10, 20]$  and  $[\min_{\text{volume}}^a, \max_{\text{volume}}^a] = [1, 5]$ . Also assume that  $a$  views the price as more important than the volume by assigning a higher weight to price, where  $[w_{\text{price}}^a, w_{\text{volume}}^a] = [0.8, 0.2]$ . Finally, let the value of an offer  $x$ , for an issue  $j$ ,  $V_j^a(x_j)$ , be modelled as a linear function:

$$V_{\text{price}}^a(x_{\text{price}}) = \frac{x_{\text{price}} - \min_{\text{price}}^a}{\max_{\text{price}}^a - \min_{\text{price}}^a},$$

$$V_{\text{volume}}^a(x_{\text{volume}}) = 1 - \frac{x_{\text{volume}} - \min_{\text{volume}}^a}{\max_{\text{volume}}^a - \min_{\text{volume}}^a}.$$

Now consider two contracts,  $[11, 5]$  and  $[15, 2]$ , offered by a client  $b$  to the server  $a$ . Given the above parameters for  $a$ , the value for the first offered price by  $b$  is  $(11 - 10/20 - 10) = 0.1$ , while the value for the first requested volume is  $(1 - (5 - 1/5 - 1)) = 0$ . The *total* value for the offered contract is the sum of the weighted values for each individual issue (namely,  $0.8 * 0.1 + 0.2 * 0 = 0.08$ ). By the same reasoning, the value of the second contract from  $b$ , on the other hand, is 0.55. Since rational action is to maximise its utility,  $a$  therefore chooses the second contract offered by  $b$  and rejects the first.

<sup>5</sup> Nonlinear approaches to modelling utility could be used if necessary, without affecting the basic ideas of the model.

<sup>6</sup> Any contract not on this frontier is sub-optimal (i.e. not Pareto-optimal) in that possible mutual gains are missed.

### 3.2. Service-oriented negotiation requirements

The above bilateral negotiation model may be valid for some service-oriented settings. However, the model contains several implicit assumptions that, although they permit good optimisation results, are inappropriate for our needs:

- (i) *Privacy of information.* To find the optimum value, the scoring functions have to be disclosed. This is, in general, inappropriate for competitive negotiation.
- (ii) *Privacy of models.* Both negotiators have to use the same additive scoring model. However, the models used to evaluate offers and generate counter offers are one of the things that negotiators try to hide from one another.
- (iii) *Value restrictions.* There are pre-defined value regions for discussion (they are necessary to define the limits of the scoring function). However, it is impossible to find these common regions and in many cases negotiation actually involves determining whether such regions even exist.
- (iv) *Time restrictions.* There is no notion of timing issues in the negotiation. However, time is a major constraint on agents' behaviour [8]. This is mainly true on the client side; agents often have strict deadlines by when the negotiation must be completed. For instance, a video link has to be provided at 16:00 because at that time a conference should start; negotiation about set up cannot continue after that time.
- (v) *Resource restrictions.* There is no notion of resource issues in the negotiation. However, the quantity of a particular resource has a strong and direct influence on the behaviour of agents, and, moreover, the correct appreciation of the remaining resources is an essential characteristic of good negotiators. Resources from the client's point of view relate directly to the number of servers engaged in the ongoing negotiation; likewise from the server's point of view. Thus, the quantity of resource has a similar effect on the agents' behaviour as time.

Even just taking the first consideration alone, it is clear that optimal solutions cannot be found in our domains: it is not possible to optimize an unknown function. Hence, we shall propose a model for individual agent negotiation that seeks to find deals acceptable to its acquaintances but which, nevertheless, maximises the agent's own scoring function.

### 3.3. A service-oriented negotiation model

In service-oriented negotiations, agents can undertake two possible roles that are, in principle, in conflict. Hence we shall distinguish (for notational convenience) two subsets of agents,<sup>7</sup> Agents = Clients  $\cup$  Servers. We use roman letters to represent agents;  $c, c_1, c_2, \dots$  will stand for clients,  $s, s_1, s_2, \dots$  for servers and  $a, a_1, b, d, e, \dots$  for non-specific agents.

We adhere to an additive scoring system (Section 3.1) in which, for simplicity, the function  $V_j^a$  is either monotonically increasing or monotonically decreasing.

In general, clients and servers have opposing interests, e.g. a client wants a low price for a service, whereas the potential servers attempt to obtain the highest price. High quality is desired by clients, but not by servers, and so on. Therefore, in the space of negotiation values, negotiators represent opposing forces in each one of the dimensions. In consequence, the scoring functions verify that given a client  $c$  and a server  $s$  negotiating values for issue  $j$ , then if  $x_j, y_j \in [\min_j, \max_j]$  and  $x_j \geq y_j$  then  $(V_j^c(x_j) \geq V_j^c(y_j) \text{ iff } V_j^s(x_j) \leq V_j^s(y_j))$ . However, in a small number of cases the clients and service providers may have a mutual interest for a negotiation issue. For example, Raiffa cites a case [12, pp. 133–147] in which the Police Officers Union and the City Hall realise, in the course of their negotiations, that they both want the police commissioner fired. Having recognised this mutual interest, they quickly agree that this course of action should be selected. Thus, in general, where there is a mutual interest, the variable will be assigned one of its extreme values. Hence these variables can be removed from the negotiation set. For instance, the act of firing the police commissioner can be removed from the set of issues under negotiation and assigned the extreme value "done".

<sup>7</sup> The subsets are not disjoint; an agent can participate as a client in one negotiation and as a service provider in another.

Once the agents have determined the set of variables over which they will negotiate, the negotiation process between two agents ( $a, b \in \text{Agents}$ ) consists of an alternate succession of offers and counter offers of values for these variables. This continues until an offer or counter offer is accepted by the other side or one of the partners terminates negotiation (e.g. because the time deadline is reached without an agreement being in place). Negotiation can be initiated by clients or servers.

We represent by  $x_{a \rightarrow b}^t$  the vector of values proposed by agent  $a$  to agent  $b$  at time  $t$ , and by  $x_{a \rightarrow b}^t[j]$  the value for issue  $j$  proposed from  $a$  to  $b$  at time  $t$ . The range of values acceptable to agent  $a$  for issue  $j$  will be represented as the interval  $[\min_j^a, \max_j^a]$ . For convenience, we assume a common global time (the calendar time) represented by a linearly ordered set of instants, namely Time, and a reliable communication medium introducing no delays in message transmission (so we can assume that transmission and reception times are identical). The common time assumption is not too strong for our application domains, because time granularity and offer and counter offers frequencies are not high. Then we have the following definition.

**Definition 1.** A *Negotiation Thread* between agents  $a, b \in \text{Agents}$ , at time  $t_n \in \text{Time}$ , noted  $X_{a \leftrightarrow b}^{t_n}$ , is any finite sequence of length  $n$  of the form  $(x_{a \rightarrow b}^{t_1}, x_{b \rightarrow a}^{t_2}, x_{a \rightarrow b}^{t_3}, \dots)$  with  $t_1, t_2 \leq \dots \leq t_n$ , where:

- (i)  $t_{i+1} > t_i$ , the sequence is ordered over time,
- (ii) For each issue  $j$ ,  $x_{a \rightarrow b}^i[j] \in [\min_j^a, \max_j^a]$ ,  $x_{b \rightarrow a}^{i+1}[j] \in [\min_j^b, \max_j^b]$  with  $i = 1, 3, 5, \dots$ , and optionally the last element of the sequence is one of the particles {accept, reject}.

We say a negotiation thread is *active*<sup>8</sup> if  $\text{last}(X_{a \leftrightarrow b}^{t_n}) \notin \{\text{accept, reject}\}$ , where last is a function returning the last element in a sequence.

For notational simplicity, we assume that  $t_1$  corresponds to the initial time value, that is  $t_1 = 0$ . In other words, there is a local time for each negotiation thread, that starts with the utterance of the first offer. When agent  $a$  receives an offer from agent  $b$  at time  $t$ ,  $x_{b \rightarrow a}^t$ , it has to rate the offer using its scoring function. If the value of  $V^a(x_{b \rightarrow a}^t)$  is greater than the value of the counter offer agent  $a$  is ready to send at the time  $t'$  when the evaluation is performed, that is  $x_{a \rightarrow b}^{t'}$  with  $t' > t$ , then agent  $a$  accepts. Otherwise, the counter offer is submitted. Expressing this concept more formally:

**Definition 2.** Given an agent  $a$  and its associated scoring function  $V^a$ ,  $a$ 's *interpretation* at time  $t'$  of an offer  $x_{b \rightarrow a}^t$  sent at time  $t < t'$ , is defined as

$$I^a(t', x_{b \rightarrow a}^t) = \begin{cases} \text{reject} & \text{if } t' > t_{\max}^a, \\ \text{accept} & \text{if } V^a(x_{b \rightarrow a}^t) \geq V^a(x_{a \rightarrow b}^{t'}), \\ x_{a \rightarrow b}^{t'} & \text{otherwise.} \end{cases}$$

where  $x_{a \rightarrow b}^{t'}$  is the contract that agent  $a$  would offer to  $b$  at the time of the interpretation, and  $t_{\max}^a$  is a constant that represents the time by which  $a$  must have completed the negotiation.

The result of  $I^a(t', x_{b \rightarrow a}^t)$  is used to extend the current negotiation thread between the agents. This interpretation formulation also allows us to model the fact that a contract unacceptable today can be accepted tomorrow merely by the fact that time has passed.

In order to prepare a counter offer,  $x_{a \rightarrow b}^{t'}$ , agent  $a$  uses a set of tactics that generate new values for each variable in the negotiation set. Based on the needs of our business process applications (Section 2), we developed the following families of tactics:

<sup>8</sup> We assume that any offer is valid (that is, the agent that uttered it is committed) until a counter offer is received. If the response time is relevant, it can be included in the set of issues under negotiation.

- (i) *Time dependent*. If an agent has a time deadline by which an agreement must be in place, these tactics model the fact that the agent is likely to concede more rapidly as the deadline approaches. The shape of the curve of concession, a function depending on time, is what differentiates tactics in this set.
- (ii) *Resource dependent*. These tactics model the pressure in reaching an agreement that the limited resources – e.g. remaining bandwidth to be allocated, money, or any other – and the environment – e.g. number of clients, number of servers or economic parameters – impose upon the agent's behaviour. The functions in this set are similar to the time dependent functions except that the domain of the function is the quantity of resources available instead of the remaining time.
- (iii) *Behaviour dependent or Imitative*. In situations in which the agent is not under a great deal of pressure to reach an agreement, it may choose to use imitative tactics to protect itself from being exploited by other agents. In this case, the counter offer depends on the behaviour of the negotiation opponent. The tactics in this family differ in which aspect of their opponent's behaviour they imitate, and to what degree the opponent's behaviour is imitated.

We do not claim that these family types are complete, nor that we have enumerated all possible instances of tactics within a given family. Rather these are merely the types of tactics we found useful in our applications.

#### 4. Negotiation tactics

*Tactics* are the set of functions that determine how to compute the value of an issue (price, volume, duration, quality, etc.), by considering a *single* criterion (time, resources, etc.). The set of values for the negotiation issue are then the range of the function, and the single criterion is its domain. The criteria we have chosen, as explained in the previous section, are time, resources and previous offers and counter offers.

Given that agents may want to consider more than one criterion to compute the value for a single issue, we model the generation of counter proposals as a weighted combination of different tactics covering the set of criteria. The values so computed for the different issues will be the elements of the counter proposal.<sup>9</sup> For instance, if an agent wants to counter propose taking into account two criteria: the remaining time and the previous behaviour of the opponent, it can select two tactics: one from the time dependent family and one from the imitative family. Both of these tactics will suggest a value to counter propose for the issue under negotiation. The actual value which is counter proposed will be the weighted combination of the two independently generated values.

To illustrate these points consider the following example. Given an issue  $j$ , for which a value is under negotiation, an agent  $a$ 's initial offer corresponds to a value in the issue's acceptable region (i.e. in  $[\min_j^a, \max_j^a]$ ). For instance, if  $a$ 's range is  $[\text{£}0, \text{£}20]$  for the price  $p$  to pay for a good, then it may start the negotiation process by offering the server  $\text{£}10$  – what initial offer should be chosen is something the agent can learn by experience. The server, agent  $b$ , with range  $[\text{£}17, \text{£}35]$  may then make an initial counter-offer of  $\text{£}25$ . With these two initial values, the strategy of agent  $a$  may consist of using a (single criterion) time dependent tactic which might make a reasonably large concession and suggest  $\text{£}15$  since it does not have much time in which to reach an agreement. Agent  $b$ , on the other hand, may choose to use two criteria to compute its counterproposal – e.g. a time dependent tactic (which might suggest a small concession to  $\text{£}24$  since it has a long time until the deadline) and an imitative tactic (which might suggest a value of  $\text{£}20$  to mirror the  $\text{£}5$  shift of the opponent). If agent  $b$  rates the time dependent behaviour three times as important as the imitative behaviour, then the value of the counter-offer will be  $(0.75 * 24) + (0.25 * 20) = \text{£}23$ . This process continues until the agents converge on a mutually acceptable solution. The origin, and subsequent evolution of, these relative weightings may be the result of expert domain knowledge, experience derived from previous negotiation cases or conditional on other factors. This is an issue which is the subject of future research.

<sup>9</sup> Values for different issues may be computed by different weighted combinations of tactics.

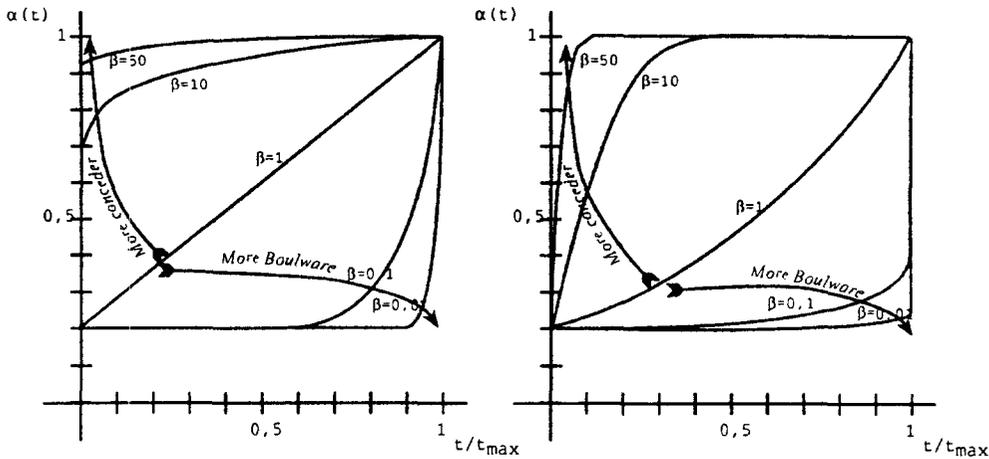


Fig. 2. Polynomial (left) and exponential (right) functions for the computation of  $\alpha(t)$ . Time is presented as relative to  $t_{\max}^a$ .

It should be noted that not all tactics can be applied at all instants. For instance, a tactic that imitates the behaviour of an opponent is only applicable when the opponent has shown its behaviour sufficiently. For this reason, the following description of the tactics pays particular attention to their applicability conditions.

#### 4.1. Time dependent tactics

In these tactics, the predominant factor used to decide which value to offer next is time,  $t$ . Thus these tactics consist of varying the acceptance value for the issue depending on the remaining negotiation time (an important requirement in our domain – Section 2.2), modelled as the above defined constant  $t_{\max}^a$ . We model the initial offer as being a point in the interval of values of the issue under negotiation. Hence, agents define a constant  $\kappa_j^a$  that when multiplied by the size of the interval, determines the value of issue  $j$  to be offered in the first proposal by agent  $a$ .

We model the value to be uttered by agent  $a$  to agent  $b$  for issue  $j$  as the offer at time  $t$ , with  $0 \leq t \leq t_{\max}^a$ , by a function  $\alpha_j^a$  depending on time as the following expression shows:

$$x_{a \rightarrow b}^t[j] = \begin{cases} \min_j^a + \alpha_j^a(t)(\max_j^a - \min_j^a) & \text{if } V_j^a \text{ is decreasing,} \\ \min_j^a + (1 - \alpha_j^a(t))(\max_j^a - \min_j^a) & \text{if } V_j^a \text{ is increasing.} \end{cases}$$

A wide range of time dependent functions can be defined simply by varying the way in which  $\alpha_j^a(t)$  is computed. However, functions must ensure that  $0 \leq \alpha_j^a(t) \leq 1$ ,  $\alpha_j^a(0) = \kappa_j^a$  and  $\alpha_j^a(t_{\max}^a) = 1$ . That is, the offer will always be between the value range, at the beginning it will give the initial constant and when the time deadline is reached the tactic will suggest to offer the reservation value.<sup>10</sup> We distinguish two families of functions with this intended behaviour: polynomial and exponential (naturally, others could also be defined). Both families are parametrised by a value  $\beta \in \mathbb{R}^+$  that determines the convexity degree (see Fig. 2) of the curve. We chose these two families of functions because of the very different way they model concession. For the same large value of  $\beta$ , the polynomial function concedes faster at the beginning than the exponential one, then they behave

<sup>10</sup> The reservation value for issue  $j$  of agent  $a$  represents the value that gives the smallest score for function  $V_j^a$ . The reservation value for agent  $a$  and issue  $j$  depends on the function  $V_j^a$  and the range  $[\min_j^a, \max_j^a]$ . If  $V_j^a$  is monotonically increasing, then the value is  $\min_j^a$ ; if it is decreasing, the reservation value is  $\max_j^a$ .

similarly. For a small value of  $\beta$ , the exponential function waits longer than the polynomial one before it starts conceding:

- *Polynomial.*  $\alpha_j^a(t) = \kappa_j^a + (1 - \kappa_j^a)(\min(t, t_{\max}^a)/t_{\max}^a)^{1/\beta}$ ,
- *Exponential.*  $\alpha_j^a(t) = \exp^{(1 - \min(t, t_{\max}^a)/t_{\max}^a)^\beta} \ln \kappa_j^a$ .

These families of functions represent an infinite number of possible tactics, one for each value of  $\beta$ . However, to better understand their behaviour we have classified them, depending on the value of  $\beta$ , into two extreme sets showing clearly different patterns of behaviour. Other sets in between these two could also be defined:

- (i) *Boulware tactics* [12, p. 48]. Either exponential or polynomial functions with  $\beta < 1$ . This tactic maintains the offered value until the time is almost exhausted, whereupon it concedes up to the reservation value.<sup>11</sup> The behaviour of this family of tactics with respect to  $\beta$  is easily understood taking into account that

$$\lim_{\beta \rightarrow 0^+} \exp\left(1 - \frac{\min(t, t_{\max}^a)}{t_{\max}^a}\right)^\beta \ln \kappa_j^a = \kappa_j^a \quad \text{or} \quad \lim_{\beta \rightarrow 0^+} \kappa_j^a + (1 - \kappa_j^a) \left(\frac{\min(t, t_{\max}^a)}{t_{\max}^a}\right)^{1/\beta} = \kappa_j^a.$$

- (ii) *Conceder* [11, p. 20]. Either exponential or polynomial functions with  $\beta > 1$ . The agent quickly goes to its reservation value. For similar reasons as before, we have

$$\lim_{\beta \rightarrow +\infty} \exp\left(1 - \frac{\min(t, t_{\max}^a)}{t_{\max}^a}\right)^\beta \ln \kappa_j^a = 1 \quad \text{or} \quad \lim_{\beta \rightarrow +\infty} \kappa_j^a + (1 - \kappa_j^a) \left(\frac{\min(t, t_{\max}^a)}{t_{\max}^a}\right)^{1/\beta} = 1.$$

## 4.2. Resource dependent tactics

These tactics are similar to the time dependent ones. Indeed, time dependent tactics can be seen as a type of resource dependent tactic in which the sole resource considered is time. Whereas time vanishes constantly up to its end, other resources may have different patterns of usage. We model resource dependent tactics in the same way as time dependent ones, that is, by using the same functions, but by either: (i) making the value of  $t_{\max}^a$  dynamic (Section 4.2.1), or (ii) making the function  $\alpha$  depend on an estimation of the amount of a particular resource (Section 4.2.2).

### 4.2.1. Dynamic deadline tactics

The dynamic value of  $t_{\max}^a$  represents a heuristic about how many resources are in the environment. The scarcer the resource, the more urgent the need for an agreement. In our application domains, the most important resource to model is the number of agents negotiating with a given agent and how keen they are to reach agreements. On one hand, the greater the number of agents who are negotiating with agent  $a$  for a particular service  $s$ , the lower the pressure on  $a$  to reach an agreement with any specific individual. While on the other hand, the longer the negotiation thread, the greater the pressure on  $a$  to come to an agreement. Hence, representing the set of agents negotiating with agent  $a$  at time  $t$  as:  $N^a(t) = \{i \mid X_{i \leftrightarrow a}^t \text{ is active}\}$ , we define the dynamic time deadline for agent  $a$  as

$$t_{\max}^a = \mu^a \frac{|N^a(t)|^2}{\sum_i |X_{i \leftrightarrow a}^t|},$$

where  $\mu^a$  represents the time agent  $a$  considers reasonable to negotiate with a single agent and  $|X_{i \leftrightarrow a}^t|$  represents the length of the current thread between  $i$  and  $a$ . Notice that the number of agents is in the numerator, so quantity of time is directly proportional to it, and averaged length of negotiation thread is in the denominator, so quantity of time is inversely proportional to it.

<sup>11</sup> Besides the pattern of concession that these functions model, Boulware negotiation tactics presume that the interval of values for negotiation is narrow. Hence, when the deadline is reached and  $\alpha(t_{\max}^a) = 1$ , the offer generated is not substantially different from the initial one.

#### 4.2.2. Resource estimation tactics

These tactics generate counter-offers depending on how a particular resource is being consumed. Resources could be money being transferred among agents, the number of agents interested in a particular negotiation, and also, in a similar way as before, time. We want the agent to become progressively more conciliatory as the quantity of resource diminishes. The limit when the quantity of the resource approaches nil is to concede up to the reservation value for the issue(s) under negotiation. When there is plenty of resource, a more Boulware behaviour is to be expected. Formally, this can be modelled by having a different computation for the function  $\alpha$ :

$$\alpha_j^a(t) = \kappa_j^a + (1 - \kappa_j^a) e^{-\text{resource}^a(t)}$$

where the function  $\text{resource}^a(t)$  measures the quantity of the resource at time  $t$  for agent  $a$ . Examples of functions are: (i)  $\text{resource}^a(t) = |N^a(t)|$ , (ii)  $\text{resource}^a(t) = \mu^a |N^a(t)|^2 / \sum_i |X_{i \leftrightarrow a}^t|$  and (iii)  $\text{resource}^a(t) = \min(0, t_{\max}^a - t)$ . In the first example, the number of negotiating agents is the resource. That is, the more agents negotiating the less pressure to make concessions. The second example models time as a resource in a similar way as in the previous section. The more agents, the less pressure, and the longer the negotiations the more pressure. Finally, the last case also models time as a resource, but in this case the quantity of resource decreases in a linear fashion with respect to time.

#### 4.3. Behaviour dependent tactics

This family of tactics compute the next offer based on the previous attitude of the negotiation opponent. These tactics have proved important in co-operative problem-solving negotiation settings [1], and so are useful in a subset of our contexts (see Section 2.2). The main difference between the tactics in this family is in the type of imitation they perform. One family imitates proportionally, another in absolute terms, and the last one computes the average of the proportions in a number of previous offers. Hence, given a negotiation thread  $\{\dots, x_{b \rightarrow a}^{t_{n-2\delta}}, x_{a \rightarrow b}^{t_{n-2\delta+1}}, x_{b \rightarrow a}^{t_{n-2\delta+2}}, \dots, x_{b \rightarrow a}^{t_{n-2}}, x_{a \rightarrow b}^{t_{n-1}}, x_{b \rightarrow a}^{t_n}\}$ , with  $\delta \geq 1$ , we distinguish the following families of tactics:

- (i) *Relative Tit-For-Tat*. The agent reproduces, in percentage terms, the behaviour that its opponent performed  $\delta \geq 1$  steps ago. The condition of applicability of this tactic is  $n > 2\delta$ .

$$x_{a \rightarrow b}^{t_{n+1}}[j] = \min \left( \max \left( \frac{x_{b \rightarrow a}^{t_{n-2\delta}}[j]}{x_{b \rightarrow a}^{t_{n-2\delta+2}}[j]} x_{a \rightarrow b}^{t_{n-1}}[j], \min_j^a \right), \max_j^a \right).$$

- (ii) *Random Absolute Tit-For-Tat*. The same as before but in absolute terms. This means that if the other agent decreases its offer by £2, then the next response should be increased by the same £2. Moreover, we add a component that modifies that behaviour by increasing or decreasing (depending on the value of parameter  $s$ ) the value of the answer by a random amount. (This is introduced as it can enable the agents to escape from local minima.)  $M$  is the maximum amount by which an agent can change its imitative behaviour. The condition of applicability is again  $n > 2\delta$ .

$$x_{a \rightarrow b}^{t_{n+1}}[j] = \min \left( \max(x_{a \rightarrow b}^{t_{n-1}}[j] + (x_{b \rightarrow a}^{t_{n-2\delta}}[j] - x_{b \rightarrow a}^{t_{n-2\delta+2}}[j]) + (-1)^s R(M), \min_j^a), \max_j^a \right)$$

where

$$s = \begin{cases} 0 & \text{if } V_j^a \text{ is decreasing} \\ 1 & \text{if } V_j^a \text{ is increasing.} \end{cases}$$

and  $R(M)$  is a function that generates a random integer in the interval  $[0, M]$ .

(iii) *Averaged Tit-For-Tat*. The agent computes the average of percentages of changes in a window of size  $\gamma \geq 1$  of its opponents history when determining its new offer. When  $\gamma = 1$  we have the relative Tit-For-Tat tactic with  $\delta = 1$ . The condition of applicability for this tactic is  $n > 2\gamma$ .

$$x_{a \rightarrow b}^{t_{n+1}}[j] = \min \left( \max \left( \frac{x_{b \rightarrow a}^{t_{n-2\gamma}}[j]}{x_{b \rightarrow a}^{t_n}[j]} x_{a \rightarrow b}^{t_{n-1}}[j], \min_j^a \right), \max_j^a \right).$$

## 5. Negotiation strategies

The aim of agent  $a$ 's negotiation strategy is to determine the best course of action which will result in an agreement on a contract  $x$  that maximises its scoring function  $V^a$ . In practical terms, this equates to how to prepare a new counter offer.

In our model, we consider that the agent has a representation of its mental state containing information about its beliefs, its knowledge of the environment (time, resources, etc.), and any other attitudes (desires, goals, obligations, intentions, etc.) the agent designer considers appropriate.<sup>12</sup> The mental state of agent  $a$  at time  $t$  is noted as  $MS_a^t$ . We denote the set of all possible mental states for agent  $a$  as  $MS_a$ .

When agent  $a$  receives an offer from agent  $b$ , it becomes the last element in the current negotiation thread between both agents. If the offer is unsatisfactory, agent  $a$  generates a counter offer. As discussed earlier, different combinations of tactics can be used to generate counter offers for particular issues. An agent's strategy determines which combination of tactics should be used at any one instant.

**Definition 3.** Given a negotiation thread between agents  $a$  and  $b$  at time  $t_n$ ,  $X_{a \leftrightarrow b}^{t_n}$ , over domain  $D = D_1 \times \dots \times D_p$ , with  $\text{last}(X_{a \leftrightarrow b}^{t_n}) = x_{b \rightarrow a}^{t_n}$ , and a finite set of  $m$  tactics.<sup>13</sup>  $T^a = \{\tau_i \mid \tau_i : MS_a \rightarrow D\}_{i \in [1, m]}$ , a *weighted counter proposal*,  $x_{a \rightarrow b}^{t_{n+1}}$ , is a linear combination of the tactics given by a matrix of weights  $\Gamma_{a \rightarrow b}^{t_{n+1}}$

$$\Gamma_{a \rightarrow b}^{t_{n+1}} = \begin{pmatrix} \gamma_{11} & \gamma_{12} & \dots & \gamma_{1m} \\ \gamma_{21} & \gamma_{22} & \dots & \gamma_{2m} \\ \vdots & \vdots & & \vdots \\ \gamma_{p1} & \gamma_{p2} & \dots & \gamma_{pm} \end{pmatrix}$$

defined in the following way:

$$x_{a \rightarrow b}^{t_{n+1}}[j] = (\Gamma_{a \rightarrow b}^{t_{n+1}} * T^a(MS_a^{t_{n+1}}))[j, j],$$

where  $(T^a(MS_a^{t_{n+1}}))[i, j] = (\tau_i(MS_a^{t_{n+1}}))[j]$ ,  $\gamma_{ji} \in [0, 1]$  and for all issues  $j$ ,  $\sum_{i=1}^m \gamma_{ji} = 1$ .

The weighted counter proposal extends the current negotiation thread as follows ( $\bullet$  is the sequence concatenation operation):

$$X_{a \leftrightarrow b}^{t_{n+1}} = X_{a \leftrightarrow b}^{t_n} \bullet x_{a \rightarrow b}^{t_{n+1}}.$$

An example of when this weighted combination is useful is when modelling a *smooth* transition from a behaviour based on a single tactic (e.g. Boulware, because the agent has plenty of time to reach an agreement) to another one (e.g. Conceder, because time is running out). Smoothness is obtained by changing the weight affecting the tactics progressively (e.g. from 1 to 0 and from 0 to 1 in the example).

<sup>12</sup> We do not prescribe a particular mental state, but rather aim towards an architecturally neutral description to ensure maximum generality for our model.

<sup>13</sup> This definition uses the natural extension of tactics to the multi-dimensional space of issues' values.

We model many-parties negotiations by means of a set of interacting negotiation threads. The way this is done is by making a negotiation thread influence the selection of which matrix  $\Gamma$  is to be used in other negotiation threads. Thus, we have:

**Definition 4.** Given  $a, b \in \text{Agents}$ , a *negotiation strategy* for agent  $a$  is any function  $f$  such that, given  $a$ 's mental state at time  $t_n$ ,  $\text{MS}_a^{t_n}$ , and a matrix of weights at time  $t_n$ ,  $\Gamma_{a \rightarrow b}^{t_n}$ , generates a new matrix of weights for time  $t_{n+1}$ , i.e.

$$\Gamma_{a \rightarrow b}^{t_{n+1}} = f(\Gamma_{a \rightarrow b}^{t_n}, \text{MS}_a^{t_n}).$$

A simplistic example of the application of our model would be to have a matrix  $\Gamma$  built up of 0's and 1's and having  $\Gamma_{a \rightarrow b}^{t+1} = \Gamma_{a \rightarrow b}^t$  for all  $t$ . This would correspond to using a fixed single tactic for each issue at every instant in the negotiation (see [9] for more details of the evolution of strategies).

## 6. Experimental evaluation of the negotiation model

The model we have presented defines and formalises a range of negotiation behaviours. However, we cannot predict from the theoretical model alone which of these behaviours will be successful in which negotiation contexts (there are too many interrelated variables and too wide a range of situations to consider). Therefore our approach is to empirically evaluate the main parameters of the model with the final aim of determining the most successful behaviours in various types of situations. At this stage, however, our investigation is focused on determining the behaviour and inter-dependencies of the model's basic constituent elements. This analysis will then lay the foundation for subsequent experimental work. To this end, we concentrate solely on the behaviour of pure tactics (i.e. we exclude strategies that combine several tactics).

The experiments involve selecting a particular tactic, generating a range of random environments, then allowing the agent to negotiate using the chosen tactic against an opponent who employs a range of other tactics. Various experimental measures related to the negotiations are then recorded. In particular, Section 6.1 defines the experimental environments and the tactics, Section 6.2 describes the experimental measures, and Section 6.3 describes the experimental hypotheses, the procedures and discusses of the results.

### 6.1. Environments and tactics

Environments are characterised by the number of agents they contain, the issues which are being discussed, the deadlines by when agreements must be reached, and the expectations of the agents. Since there are infinitely many potential environments, we need to find a means of selecting a representative and finite subset in which we can assess an agent's negotiation performance.

To this end, the first simplification involves limiting ourselves to bilateral negotiation between a single client and server over the single issue of *price*. Given this situation, the experimental environment is uniquely defined by the following variables:  $[t_{\max}^c, t_{\max}^s, \kappa^c, \kappa^s, \min_{\text{price}}^c, \max_{\text{price}}^c, \min_{\text{price}}^s, \max_{\text{price}}^s]$ . We compute the negotiation interval (the difference between the agent's minimum and maximum values) for price using two variables:  $\theta^a$  (the length of the reservation interval for an agent  $a$ ) and  $\Phi$  (the degree of intersection between the reservation intervals of the two agents; ranging between 0 for full overlap and 0.99 for virtually no overlap). In this case, for each environment, we assigned  $\min_{\text{price}}^c = 10$ , set  $\Phi = 0$ , randomly selected  $\theta^a$  between the ranges of  $\{10, 30\}$  for both agents and computed the negotiation intervals as  $\max^c = \min^c + \theta^c$ ;  $\min^s = \theta^c \Phi + \min^c$ ;  $\max^s = \min^s + \theta^s$ .<sup>14</sup>

<sup>14</sup> Note that the server's minimum reservation value is never lower than the client's minimum. This is because we are not interested in degenerate negotiations in which offers are immediately accepted. This method of generating reservation values also means that a deal is always possible since there is always some degree of overlap.

Tactic Family	Tactic Name	Abbreviation	Tactic Ranges	Description
Time-dependent	Boulware	B	$\beta \in \{0.01, 0.2\}$	Increased rate of approach to reservation as $\beta$ increases
Time-dependent	Linear	L	$\beta = 1.0$	
Time-dependent	Conceder	C	$\beta \in \{20.0, 40.0\}$	
Resource-dependent	Impatient	IM	$\mu = 1, n = 1$	Decreasing rate of approach to reservation as $\mu$ increases
Resource-dependent	Steady	ST	$\mu \in \{1, 5\}, n = 1$	
Resource-dependent	Patient	PA	$\mu \in \{5, 10\}, n = 1$	
Behaviour-dependent	Relative tit for tat	RE	$\delta = 1$	Percentage imitation of last two offers
Behaviour-dependent	Random tit for tat	RA	$\delta = 1, m \in \{1, 3\}$	Fluctuating absolute imitation of last two offers
Behaviour-dependent	Average tit for tat	AV	$\gamma = 2$	Average imitation of last four offers

Fig. 3. Experimental tactic key.

The second simplification involves selecting a finite range of tactics since the model allows for an infinite set (e.g. the range of  $\beta$  is infinite which means there are infinitely many time dependent tactics). For analytical tractability, we divided the tactics into nine groups (see Fig. 3); three each from the time, resource and behaviour dependent families. We chose an equal number for each family to ensure that the results are not skewed by having more encounters with a particular type of tactic.

Each tactic group is then sampled for every environment since we are interested in the behaviour of tactic families rather than single, concrete tactics. For each environment  $e_k$ ,  $k$  indexes the environments, we define two matrices representing the outcomes of the client,  $\text{game}_c^{e_k}$ , and the server,  $\text{game}_s^{e_k}$ , when playing particular tactics. We index the client's tactics by the rows  $i$  and the server's by the columns  $j$ , so  $\text{game}_c^{e_k}[i, j]$  is the outcome of the client when playing tactic  $i$  against a server playing tactic  $j$ . Each tactic plays against all other tactics in each environment, hence  $1 \leq i, j \leq 9$ .

## 6.2. Experimental measures

To evaluate the effectiveness of the tactics, we consider the following measures which calibrate: (i) the intrinsic benefit of the tactic family to an agent (Section 6.2.1); (ii) the cost adjusted benefit which moderates the intrinsic benefit with some measure of the cost involved in achieving that benefit (Section 6.2.2) and (iii) the performance of the intrinsic utility relative to a game of perfect information (Section 6.2.3).

### 6.2.1. Intrinsic agent utility

The intrinsic benefit is modelled as the agent's utility for the negotiation's final outcome, in a given environment, independently of the time taken and the resources consumed [14]. This utility,  $U_a^{e_k}$ , is calculated for each agent for a price  $x$  using a linear scoring function:<sup>15</sup>

$$U_c^{e_k}(x) = \frac{\max_{\text{price}}^c - x}{\max_{\text{price}}^c - \min_{\text{price}}^c}, \quad U_s^{e_k}(x) = \frac{x - \min_{\text{price}}^s}{\max_{\text{price}}^s - \min_{\text{price}}^s}.$$

If no deal is made in a particular negotiation, then we assign zero to both  $U_c^{e_k}$  and  $U_s^{e_k}$ . However, by defining the utilities in this manner we cannot distinguish between deals made at reservations and no deals. Therefore in certain experiments we compute the intrinsic utility only for cases in which deals are made.

The outcome of the negotiations, as presented in Section 6.1, is represented in the matrix  $\text{game}_a^{e_k}$ . Hence the utility for a client  $c$  when negotiating using a tactic  $i$  against a server  $s$  using tactic  $j$  in environment  $e_k$  is  $U_c^{e_k}(\text{game}_c^{e_k}[i, j])$ .

<sup>15</sup> We acknowledge that this is a simple utility function, but our intention here is to investigate the properties of the model and not the utility functions per se.

### 6.2.2. Cost adjusted benefit

In addition to knowing the intrinsic utility to an agent, we are also interested in knowing the relationship between an outcome's utility and the costs involved in achieving it. Therefore the cost adjusted benefit ( $B$ ) of tactic pairs  $i$  and  $j$  in environment  $e_k$  is defined as follows:

$$B_a^{ek}[i, j] = U_a^{ek}[i, j] - C_a^{ek}[i, j].$$

To define the cost function,  $C$ , we introduce the notion of a *system*. A system in these experiments is a set of resources that can be used by the agents during their negotiations. The usage of these resources is subject to a tax  $\mathcal{T}$  which is levied on each message communicated between the agents. Therefore, the greater the communication between the agents, the greater the cost to the agents. So

$$C_c^{ek}[i, j] = C_s^{ek}[i, j] = \tanh(|X_{c_i \leftrightarrow s_j}| * \mathcal{T}).$$

where  $|X_{c_i \leftrightarrow s_j}|$  is the length of the thread at the end of negotiation between a client using tactic  $i$  and a server using tactic  $j$ ,  $\tanh$  is an increasing function that maps the real numbers into  $[0, 1]$  and  $\mathcal{T}$  determines the rate of change of  $\tanh()$ . We sampled  $\mathcal{T}$  between the ranges of  $[0.001, 0.1]$ . In short, the greater the taxation system, the more costly the communication, and the quicker the rate at which the cost rises to an agent for each message.

The system utility, on the other hand, is coarsely defined as the total number of messages in negotiation which indirectly measures the communication load the tactics incur at the agent level.

### 6.2.3. Experimental controls: the perfect information game

The aforementioned measures calibrate the relative performance of the tactics within our model. However, we also need to calibrate the tactic's performance with respect to some control conditions. Game Theory, Economics, Social Choice and Voting Theory have all proposed desirable properties and solution criteria that can be used to characterise an agent's negotiation. Typically, these properties and criteria are concerned with the influence of the individual agent on the outcome or, conversely, the influence of the outcome on the individual. Specific measures include: Pareto optimality, symmetry, fairness and individual rationality [4]. We chose to compare the outcome attained by a pair of our tactic families with that suggested by a protocol in which agents declare their true reservation prices at the first step of negotiation and then share the overlap in the declared reservation values. This choice is both fair and Pareto optimal (in that the outcome is beneficial to both agents and any deviation results in an increase in utility for one at the cost of a decrease in utility to the other [3]). For example, consider a client agent  $c$  and a server agent  $s$  having price reservation values  $[\min_{\text{price}}^c, \max_{\text{price}}^c]$  and  $[\min_{\text{price}}^s, \max_{\text{price}}^s]$  respectively and  $\max_{\text{price}}^c \geq \min_{\text{price}}^s$ . We then define the control outcome  $\mathcal{O}$  for a given environment  $e_k$  as

$$\mathcal{O}^{ek} = \frac{\max_{\text{price}}^c + \min_{\text{price}}^s}{2}.$$

Applying the definitions of utility presented earlier, we can compute the utility of the control game,  $U_a^{ek}(\mathcal{O}^{ek})$ , for agent  $a$ . Given this, we define the comparative performance of agents using our model with respect to the one shot protocol, as the difference between the intrinsic agent utility and the utility the agent would have received in the control protocol:

$$\text{Gain}_a^{ek}[i, j] = U_a^{ek}(\text{game}_a^{ek}[i, j]) - U_a^{ek}(\mathcal{O}^{ek}).$$

### 6.2.4. Average utilities

To produce statistically meaningful results, we average the utilities over a number of environments and sum them against all other tactics for each agent. Therefore our analysis is based on the performance of a tactic family across all other tactic families. The precise set of environments is sampled from the parameters specified in Section 6.1 and the number of environments used is 200. This ensures that the probability of the sampled mean deviating by more than 0.01 from the true mean is less than 0.05.

### 6.3. Hypotheses and results

The experiments considered here relate to two main components of the negotiation model: (i) the amount of time available to make an agreement,  $t_{\max}^a$  and (ii) the relative value of the initial offer,  $\kappa^a$ . To test the effects of varying deadlines on agreements, we classify the experiments into environments where the time to reach an agreement is large (Section 6.3.1) and those where it is small (Section 6.3.2). Likewise for initial offers; there are environments in which the initial offer is near the minimum of the agent's reservation values and those where it is near the maximum (Section 6.3.3). The reservation values were computed as described in Section 6.1 with  $\theta^c = \theta^s = 30$  and  $\Phi = 0$ . The reader is referred to Fig. 3 for the key to the experimental tactics. Each abbreviation is further postfixed by the agent's role (e.g. BC and BS denote a client and a server playing tactic B, respectively).

#### 6.3.1. Long term deadlines

Our hypotheses about the effect of long term deadlines can be stated as follows:

**Hypothesis 5.** In environments where there is plenty of time for negotiation, tactics which slowly approach their reservation values will gain higher intrinsic utilities than those which have a quicker rate of approach. However, they will make fewer deals.

**Hypothesis 6.** The utility to the system will be high when tactics have long deadlines since large numbers of offers will be exchanged. Consequently, there will be a large difference between a deal's intrinsic and cost adjusted utilities.

To evaluate these hypotheses we need to provide concrete values for the experimental variables. In this case, we define an environment with long term deadlines as one in which the values of  $t_{\max}^c$  and  $t_{\max}^s$  are sampled within thirty and sixty ticks of a discrete clock. Note that we allow  $t_{\max}^c \geq t_{\max}^s$  and  $t_{\max}^c < t_{\max}^s$ . Since high values of  $\kappa^a$  overconstrain the true behaviour of tactics, we set  $\kappa = 0.1$  for both agents. In each environment, the order of who begins the negotiation process is randomly selected.<sup>16</sup>

Consider Hypothesis 5 first. We predicted that a tactic which approaches reservations at the slowest rate (i.e., a Boulware) should attain the best deals. However, from Fig. 4(A) we observe that the most successful tactics are Linear, Patient and Steady. These tactics are characterised by the fact that they concede at a steady rate throughout the negotiation process. The next most successful group are the behaviour dependent tactics. Note, these imitative tactics never do better than other tactics; the best they do is gain equal utility to the best tactic [1]. The worst performing tactics are Conceder and Impatient, both of which rapidly approach their reservation values.

To help explain Boulware's unexpectedly poor performance, we note that these tactics make significantly fewer deals than all the other tactic families (Fig. 4(C)). Taking this into account, we examined the average intrinsic utility for only those cases in which deals are made (Fig. 4(B)). This shows that when Boulwares do make deals, they do indeed receive a high individual utility (as predicted).

We hypothesised that the reason why Boulware tactics perform poorly may be caused by the imitating responses of the behaviour dependent tactics, thereby effectively increasing the numbers of Boulwares in the population. To test this, we compared final average intrinsic utility for deals only of Boulware tactics across: (i) *all* other tactics and (ii) all other tactics *apart* from behaviour dependent tactics. We found that the success of Boulware tactics increased by 10% in the latter case.

From these observations, we conclude that our initial hypothesis does not hold because of the composition of the tactic population. We predicted that in an environment in which there is plenty of time to reach a deal, Boulware

<sup>16</sup> The initiator of a bid is randomly chosen because in our earlier experiments we found that the agent which opens the negotiation fairs better, irrespective of whether the agent is a client or a server. This is because the agent who begins the negotiation round reaches  $\alpha_{\text{price}}^a = 1$  before the other agent, hence deriving more intrinsic utility.

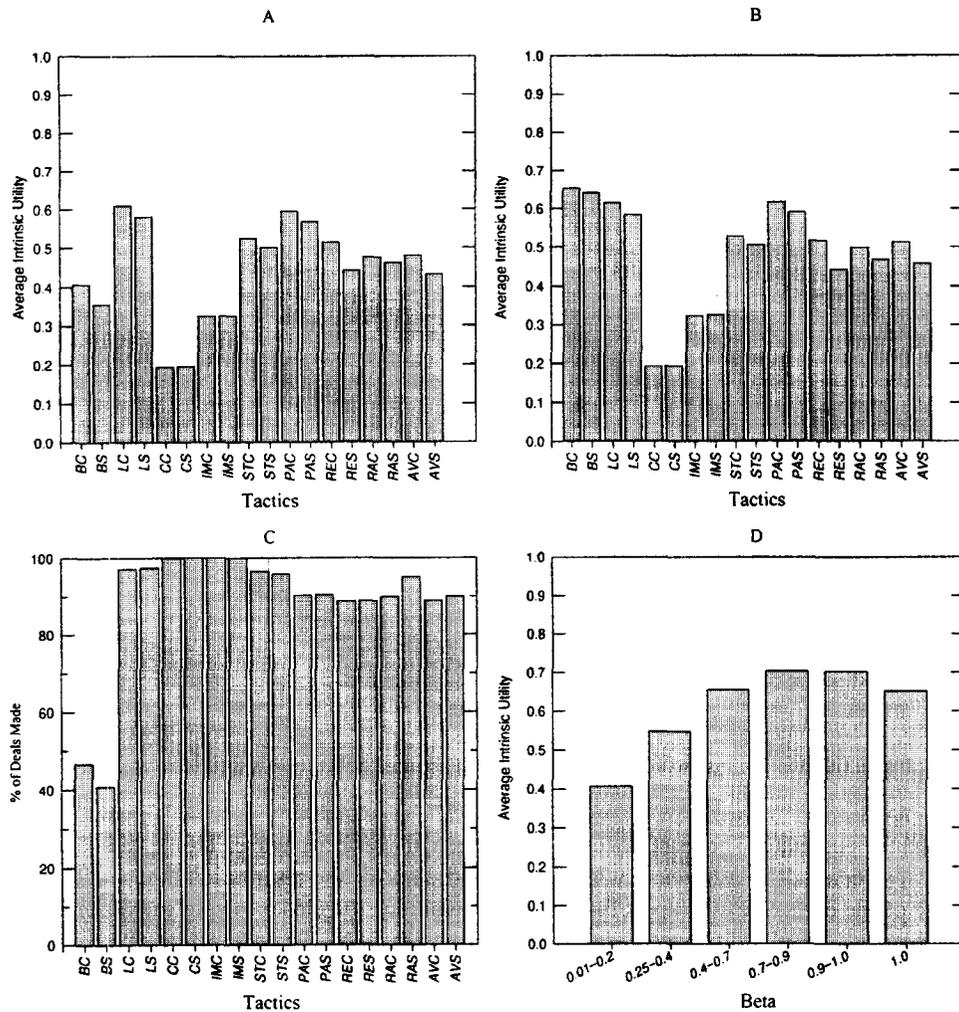


Fig. 4. (A) Average intrinsic utility for both deals and no deals. (B) Average intrinsic utility for deals only. (C) Percentage of deals made. (D) Average intrinsic utility for both deals and no deals for increasing values of  $\beta$ .

should rank higher than tactics that approached reservation values quickly. However, for Boulwares to prosper in our experimental environment, they should adopt a value for  $\beta$  which is between 0.7 and 1.0 (Fig. 4(D)).

Moving onto Hypothesis 6. Fig. 5(A) confirms the results for the first part of this hypothesis; the tactic that uses the most system resource is Boulware, and the least is Conceder. In addition, although Boulware tactics have higher intrinsic agent utilities than conciliatory tactics (Conceder and Impatient), when the the cost of communication is taken into consideration the converse is true (Fig. 5(B)). This accords with our intuitions in the second part of Hypothesis 6. The cost adjusted utilities of the remaining tactics are approximately similar. The reason for this is that cost adjusted benefit, which is the product of the intrinsic utility and a function of the number of exchanged messages, is sensitive to large fluctuations in the product and assigns similar utilities to non-extreme values.

Finally, we observe that the comparison of our tactics with respect to the controls follows the same broad pattern as the intrinsic agent utility (Fig. 5(C)). Steadily conceding type tactics (Linear, Steady and Patient) perform better

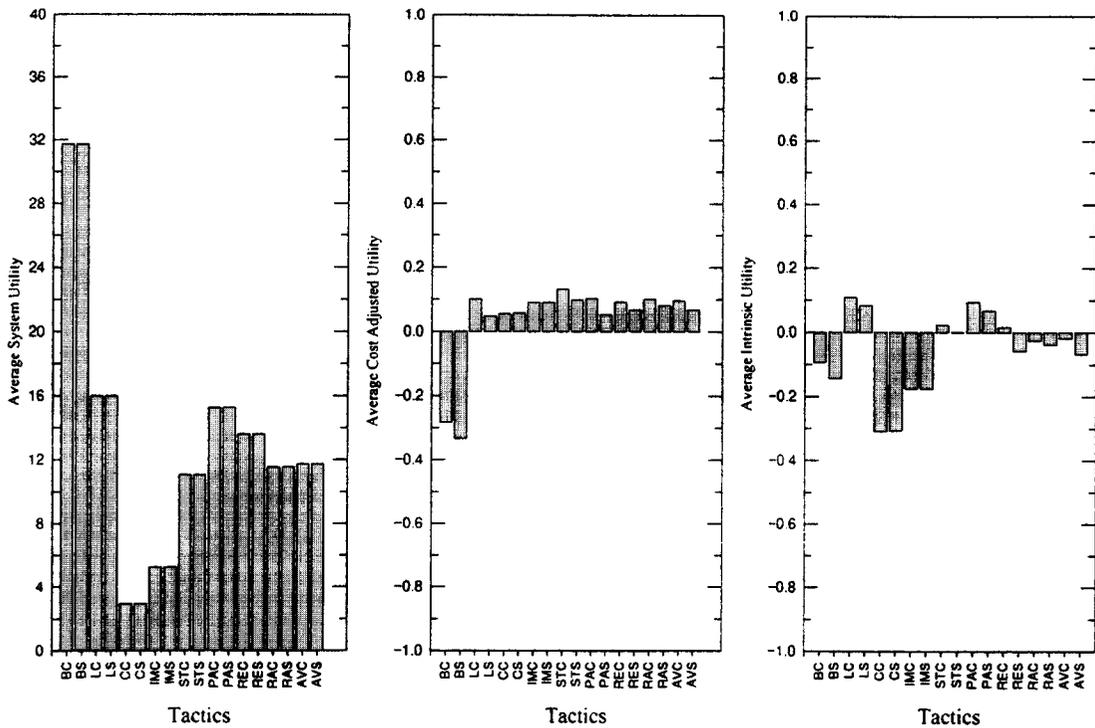


Fig. 5. (A) Average system utility. (B) Average cost adjusted utility. (C) Comparisons to control.

than the controls, the conciliatory types (Conceder and Impatient) perform worse. This is to be expected, since the closer the tactic's selected deal to the deal which is the mid-point of the reservation intersection (intrinsic utility of 0.5 – because of the complete overlap of the reservation values), the closer to zero the differential between the intrinsic utility and the control utility becomes. As we see from Fig. 4(A), the only tactics which approach or exceed an average intrinsic utility of 0.5 are those which concede at a steady rate.

### 6.3.2. Short term deadlines

Changing the environmental setting can radically alter the successfulness of a particular family of tactics. Therefore, we carried out an experiment to investigate the behaviour of tactics in cases where deadlines are short. For this case, our hypotheses are:

**Hypothesis 7.** When there is a short time frame to negotiate, tactics which quickly approach their reservation values will make more deals.

**Hypothesis 8.** Since deadlines are short, the number of messages exchanged to reach a deal will be small. Consequently the system utility will be low.

In this context, short term deadlines are obtained by sampling values for  $t_{\max}^c$  and  $t_{\max}^s$  between two and ten ticks of a discrete clock. The remainder of the experimental setup is as before.

Fig. 6 shows the results obtained for these experiments. The first observation is that for most tactics, the overall intrinsic utility, the system utility and the number of deals made (Figs. 6(A), (C) and (B), respectively) are significantly lower than the respective measures for the long deadline experiments. A lower system utility is expected

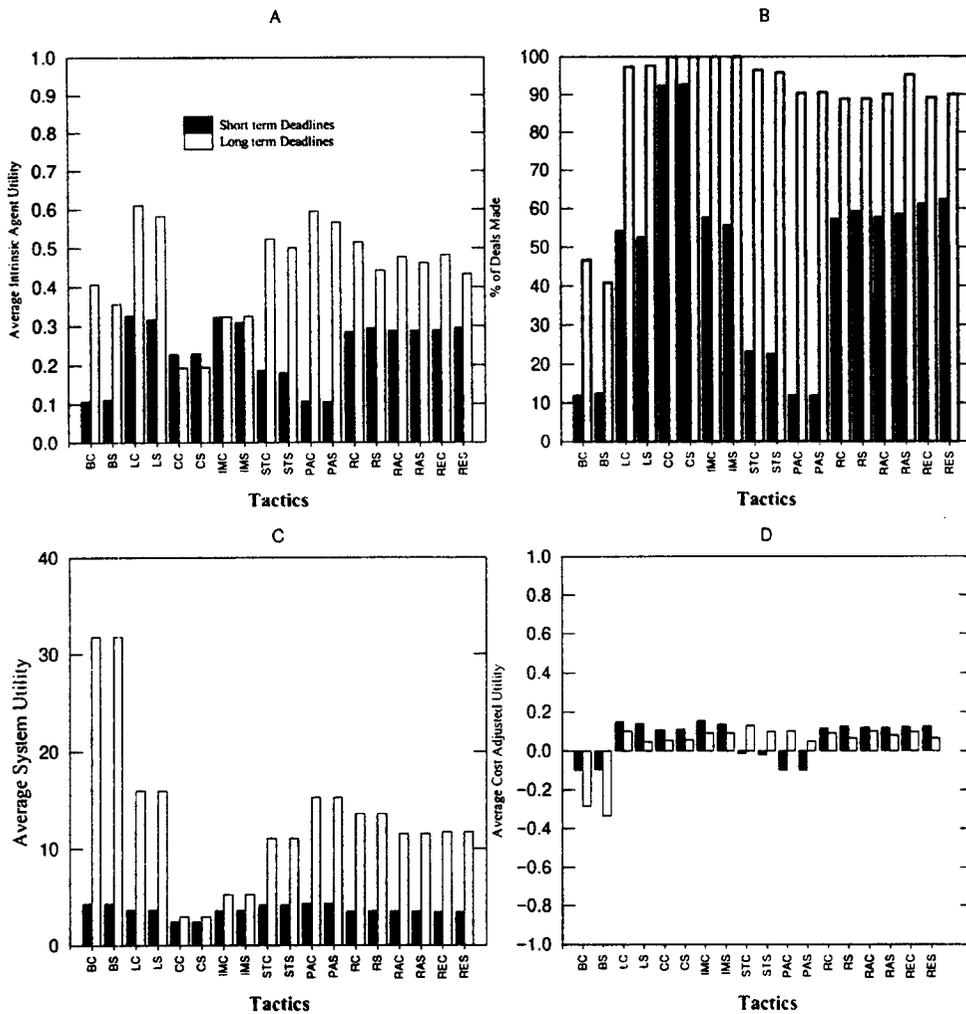


Fig. 6. Comparative data for long and short term deadlines: (A) average intrinsic utility; (B) percentage number of deals; (C) average system utility; (D) average cost adjusted utility.

since fewer messages can be exchanged in the allocated time. Note that since Conceder and Impatient are quick to reach agreements, their utilisation of system resources is independent of the time constraints. Also, because fewer messages are exchanged, the agents pay less tax and, consequently, keep a greater percentage of their derived intrinsic utility (Fig. 6(D)). These findings are all in line with the predictions in Hypothesis 8. However, the other measures require further analysis.

With long term deadlines, most tactics, apart from Boulware, make deals approximately 90–95% of the time, whereas with short term deadlines only Conceder makes anything like this number. This reduction is either because the tactics are insensitive to changes in their environment (e.g. resource dependent tactics) or because they have a slow rate of approach to reservation values (e.g. Boulware). Time insensitivity means the other tactics fail to make many deals when interacting with these tactics. Because the length of the thread is independent of the deadline, the resource dependent tactics cannot distinguish between short and long term deadlines. This claim is supported by the observation that Impatient gains equivalent intrinsic utility independently of deadlines (Fig. 6(A)). Furthermore,

resource dependent tactics are differentiated with respect to  $\mu$ , the amount of time an agent considers reasonable for negotiation. If an agent does not reason about deadlines and erroneously assumes a value for  $\mu$  which is close to or above  $t_{\max}$ , then it will be unsuccessful in environments where deadlines are important. The relatively low intrinsic utility of Patient and Steady (ranked 9th and 7th respectively – Fig. 6(A)) supports this claim. When the deadline is long, resource dependent tactics with  $\mu > 1$  gain large intrinsic utility because they approach reservation values in a steady way. However, the same behaviour in short term deadlines is less successful. The imitative tactics also exhibit a reduction in an average intrinsic utility. This is to be expected since these tactics imitate the relatively larger rate of concession of other tactics (especially time dependent tactics) when the deadline is shorter.

Hypothesis 7 is supported by the relative reductions in intrinsic utility for Boulware, Steady and Patient and by the comparative increase for Conceder and Impatient. Whereas in long term deadlines, Boulware, Steady and Patient ranked higher than the conciliatory tactics, the reverse is true for short term cases. With short term deadlines, tactics that quickly approach their reservation values gain higher intrinsic utility than those which are slower.

Again, we discovered that the dominant tactic is one which concedes at a steady rate (i.e. Linear), suggesting that the best tactic, independent of time deadlines, is one that approaches reservation values in a consistent fashion. The behaviour dependent tactics also gain relatively high utilities in both cases, ranking third and fourth for short and long term deadlines, respectively. Thus, whereas most tactics have large fluctuations in rankings across environments, the behaviour dependent family maintains a stable position, indicating its general robustness and usefulness in a wide range of contexts. This is because these tactics stick firm to avoid exploitation and reciprocate concession.

### 6.3.3. Initial offers

In our model, an agent's reservation values are private. This means no other agent has any knowledge of where in the range of acceptable values an opponent begins its bidding process, nor where it is likely to end. Given this constraint, an agent must decide where in its reservation ranges it should begin its negotiation offers. That is, what should be the value of  $\kappa^a$  in the face of this uncertainty? To help answer this question, we formed the following hypothesis:<sup>17</sup>

**Hypothesis 9.** When the deadline for agreements is not short, making initial offers which have values near the maximum of  $U_{\text{price}}^a$  leads to deals which have higher intrinsic agent utilities than initial offers near the minimum of  $U_{\text{price}}^a$ . In other words, a server that starts bidding close to  $\max_{\text{price}}^s$  is more likely to end up with deals that have a higher utility than a server who starts bidding close to  $\min_{\text{price}}^s$ . The converse is true for the client.

To test this hypothesis, we let both agents have reasonably long deadlines,  $t_{\max}^c = t_{\max}^s = 60$ , and made  $\kappa^c$  a constant at 0.1 (i.e. the client is cautious in its first offer). Therefore, the single independent variable was  $\kappa^s$ , which we sampled between the values [0.01, 0.2] for high initial price offers and [0.8, 0.99] for low initial offers. All other environmental variables were chosen as in previous experiments.

Fig. 7 confirms our prediction that a server which begins bidding at values near the maximum of  $U_{\text{price}}^s$  (Fig. 7(A)) has a higher average intrinsic utility than a server that begins bidding at values near the minimum of  $U_{\text{price}}^s$  (Fig. 7(B)). Moreover, if  $\kappa^s$  is close to  $\kappa^c$  (the client starts bidding at low values and the server begins with high offers), then both agents gain equivalent utility in most cases and take many rounds of negotiations before a deal is found (7(C)). This is because the tactics begin their negotiation at some distance from the point in the negotiation space where bids have values which have a mutually acceptable level.

Conversely, if  $\kappa^s$  is not close to  $\kappa^c$  (both the client and server start bidding at low values), then the client benefits substantially more than the server. This is because the initial offers of the server are now immediately within the

<sup>17</sup> Note:  $U_{\text{price}}^s$  increases and  $U_{\text{price}}^c$  decreases with increasing price offers.

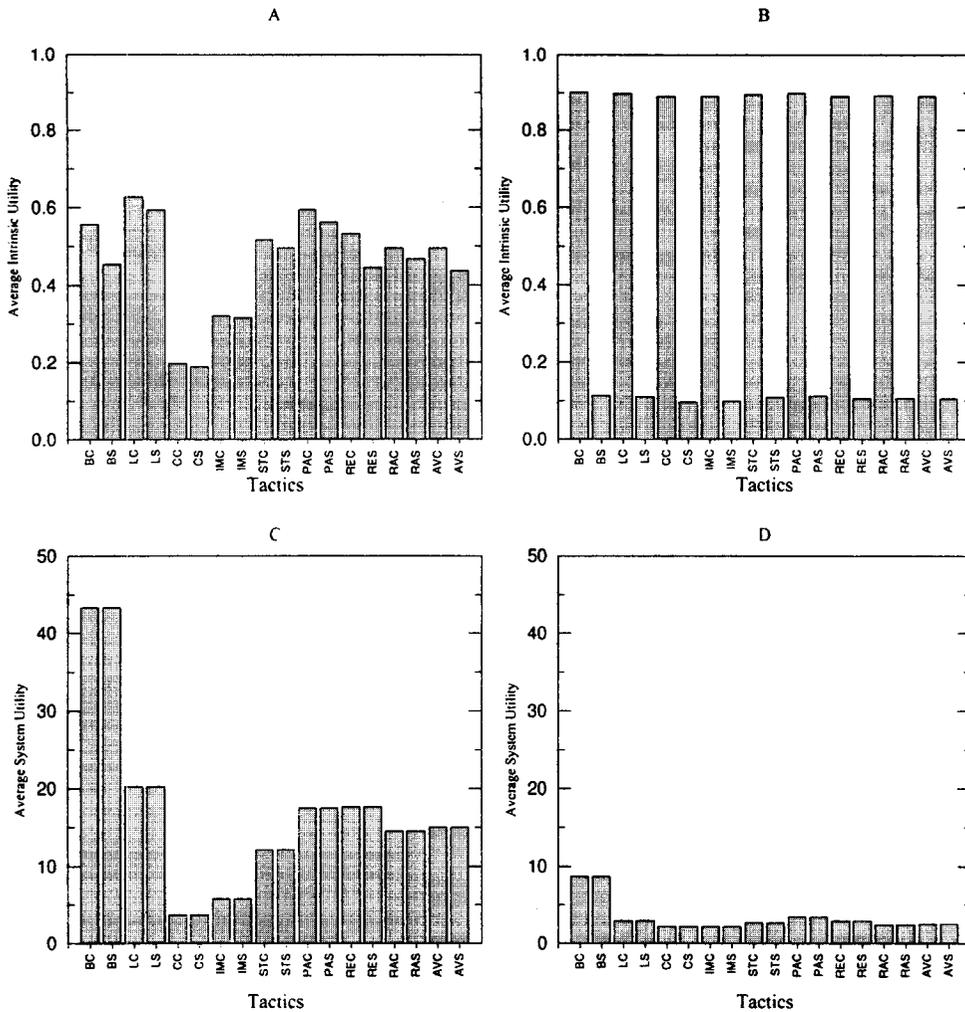


Fig. 7. (A) Average intrinsic utility for  $\kappa^s \in \{0.01, 0.2\}$ . (B) Average intrinsic utility for  $\kappa^s \in \{0.8, 0.99\}$ . (C) Average system utility for  $\kappa^s \in \{0.01, 0.2\}$ . (D) Average system utility for  $\kappa^s \in \{0.8, 0.99\}$ .  $\kappa^c = 0.1$  for all cases.

acceptance level of the client (confirmed by the number of messages exchanged before a deal is reached (Fig. 7(D)). Thus, the client gains relatively more utility than a server, since the initial offers of both agents are low and deals are made at low values.<sup>18</sup>

We can further explain the influence of  $\kappa$  on the behaviour of tactics from the observations shown in Fig. 8.  $\kappa^a$  is used by all tactics for generating the initial offer but, for exposition purposes, we only discuss the results with respect to the Boulware tactic family (since this offers the greatest difference in behaviour). When  $\kappa^s$  is low, Boulwares have a lower percentage of deals relative to other tactics (Fig. 8(A)). Conversely, when  $\kappa^s$  is high, Boulware almost equals all other tactics in the percentage of deals they make (Fig. 8(B)). This is because at low values of  $\kappa^s$ , the shape of the acceptance level for Boulware is almost a step function, whereas when  $\kappa^s$  is high it is a straight line

<sup>18</sup> When  $\kappa^s$  is distinctly different from  $\kappa^c$  there is little differentiation among intrinsic utilities. This is why we set  $\kappa^a = 0.1$  for both agents in Sections 6.3.1 and 6.3.2.

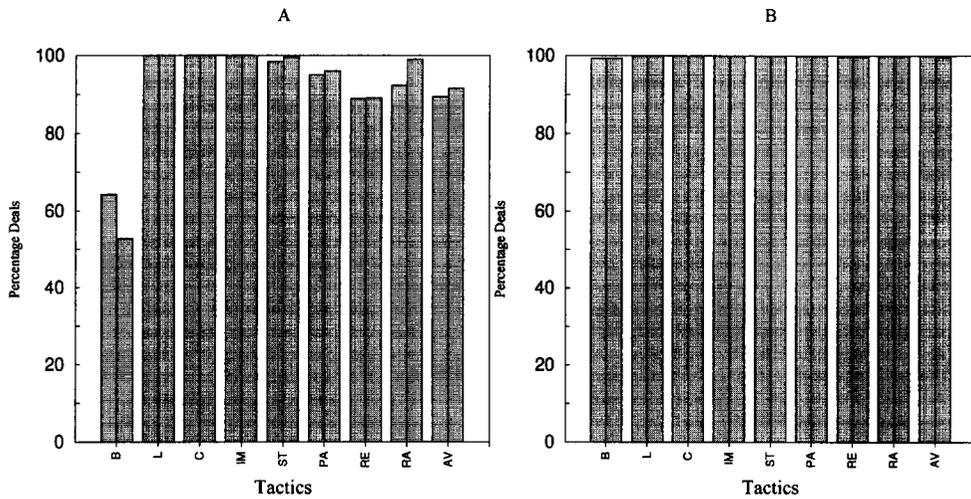


Fig. 8. Percentage of successful deals when  $\kappa^c = 0.1$  and (A)  $\kappa^s \in \{0.01, 0.2\}$ , (B)  $\kappa^s \in \{0.8, 0.99\}$ .

near to or at  $\min^s$ . Thus a server playing a Boulware tactic makes a small number of high utility deals when the acceptance levels tend towards being a step function (compare Figs. 8(A) and 7(A)), but makes larger number of lower utility deals when the acceptance level is almost a straight line (Figs. 8(B) and 7(B)). Therefore, as the value of  $\kappa$  increases, the likelihood of a deal increases, but the utility of the deal decreases.

## 7. Related work

Research in negotiation models has been pursued in different fields of knowledge: game theory, social sciences and artificial intelligence. Each field has concentrated on different aspects of negotiation, making the assumptions that were pertinent for the goal of their study. In game theory, researchers have been interested in mechanism design: the definition of protocols that limit the possible tactics (or strategies) that can be used by players. For instance, they are interested in defining protocols that give no benefit to agents that mis-represent or hide information [13]. In this work disclosure of information is acceptable, because by doing so it will benefit the agent in finding an optimal solution for itself. Contrary to our model, and as we discussed in Section 2, this is an inappropriate assumption from the point of view of real applications. As has been argued elsewhere [17], these and other assumptions limit the applicability of game theory models to solve real problems. In a paper in this issue, Wellman and Wurman present a justification of the adoption of mechanism design to situations in which disclosure of information is not possible or acceptable [16]. They present market price systems as institutions that can be used to model resource allocation in general. Our approach agrees with this point of view concerning disclosure of information, and complements it in that we concentrate more on the internal decisions of negotiating agents given a particular protocol, and not on the process of mechanism design.

Our interests lie in investigating the process of negotiation among agents and not only on the outcome. Hence, our study, and those forthcoming, are much more in the experimental line of [5]. Although we do not concentrate on learning, some similarities can be found with the formalism by Zeng and Sycara [17]. We have not concentrated however on the description of negotiation protocols that has been an important focus of attention for the community of distributed artificial intelligence (see [10] for extensive references).

## 8. Discussion and future work

This paper presented a formal model of an autonomous agent's decision function as it relates to the process of service-oriented negotiation. The model defines a number of tactics which agents can employ during negotiations and it indicates how an agent can change these tactics over time to give various forms of strategic behaviour. The form of the model, and the assumptions it makes, has been guided by our experiences in developing real-world agent applications for the domain of business process management. For this reason, the model is well suited for practical agent applications.

In earlier work [15], we proved that agents negotiating using our model were guaranteed to converge on a solution in a number of well-defined situations. In this paper, we sought to extend these results and to evaluate the model in a wider range of circumstances. To this end, we defined a number of basic hypotheses about negotiation using our model and sought to validate them empirically. In particular, with respect to tactics we discovered that: (i) irrespective of short or long term deadlines, it is best to be a linear type tactic, otherwise an imitative tactic; (ii) tactics must be responsive to changes in their environment; and (iii) there is a tradeoff between the number of deals made and the utility gained which is regulated by the initial offers.

The aforementioned results confirmed (and rebutted!) a number of basic predictions about negotiation using our model. Our aim for the future is to extend this evaluation to cover a wider range of phenomena. In particular, we intend to: (i) extend the analysis to other types of environments (for example, we predict that an increase in the number of agents will affect resource dependent tactics and dramatically influence the dynamics of all tactic interactions); (ii) investigate the effects of strategies – weighted combination of tactics may outperform pure tactics in certain environments (see [9] for some results); (iii) investigate the tactic “pool” which makes up the population (for example, we predict that the number and value of deals made between members of a society that is made up solely of Boulwares will be significantly different to societies where the population has a mixture of Boulwares and imitators). Finally, to gain further explanatory power, we intend to analyse the behaviour of tactics, in these and future environments, at the level of pairwise interactions.

## Acknowledgements

This project has received the support of the DTI/EPSRC Intelligent Systems Integration Programme (ISIP) project ADEPT, Nortel Technology and the Spanish Research project SMASH (CICYT number, TIC96-1038-C04001).

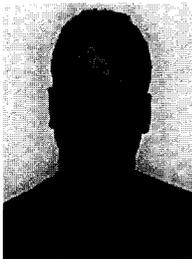
## References

- [1] R. Axelrod, *The Evolution of Cooperation*, Basic Books, New York, 1984.
- [2] B. Crabtree, N.R. Jennings (Eds.), *The Practical Application of Intelligent Agents and Multi-Agent Technology*, London, UK, 1996.
- [3] G. Debreu, *Theory of Value, An Axiomatic Analysis of Economic Equilibrium*, Wiley, New York, 1959.
- [4] E. Ephrati, J.S. Rosenschein, Multi-Agent Planning as Search for a Consensus that Maximizes Social Welfare, *Lecture Notes in Artificial Intelligence*, vol. 830, Springer, Berlin, 1994, pp. 207–226.
- [5] A. Ito, H. Yano, The emergence of cooperation in a society of autonomous agents – the prisoner's dilemma game under the disclosure of contract histories, in: V. Lesser (Ed.), *Proceedings of the First International Conference on Multi-Agent Systems*, San Francisco, CA, 1995, AAAI Press/MIT Press, Cambridge, MA, pp. 201–208.
- [6] N.R. Jennings, P. Faratin, M.J. Johnson, T.J. Norman, P. O'Brien, M.E. Wiegand, Agent-based business process management, *International Journal of Cooperative Information Systems* 5 (2–3) (1996) 105–130.
- [7] R.L. Keeney, H. Raiffa, *Decisions with Multiple Objectives*, Wiley, New York, 1976.
- [8] S. Kraus, J. Wilkenfeld, G. Zlotkin, Multiagent negotiation under time constraints, *Artificial Intelligence* 75 (2) (1995) 297–345.
- [9] N. Matos, C. Sierra, N.R. Jennings, Determining successful negotiation strategies, an evolutionary approach, *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*, Paris, France, 1998.
- [10] H.J. Mueller, Negotiation principles, in: G.M.P. O'Hare, N.R. Jennings (Eds.), *Foundations of Distributed Artificial Intelligence*, Sixth-Generation Computer Technology Series, Wiley, New York, 1996, pp. 211–229.

- [11] D.G. Pruitt, *Negotiation Behavior*, Academic Press, New York, 1981.
- [12] H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press, Cambridge, MA, 1982.
- [13] J.S. Rosenschein, G. Zlotkin, *Rules of Encounter*, MIT Press, Cambridge, MA, 1994.
- [14] S. Russell, E. Wefald, *Do the Right Thing*, MIT Press, Cambridge, MA, 1991.
- [15] C. Sierra, P. Faratin, N.R. Jennings, A service-oriented negotiation model between autonomous agents, in: M. Boman, W. Van de Velde (Eds.), *Multi-Agent Rationality, Proceedings of the Eighth European Workshop on Modelling Autonomous Agents in Multi-Agent World, MAAMAW'97, Lecture Notes in Artificial Intelligence*, vol. 1237, Springer, Berlin, 1997, pp. 17–35.
- [16] M.P. Wellman, P.R. Wurman, Market-aware agents for a multiagent world, *Robotics and Autonomous Systems* 24 (1998) this issue.
- [17] D. Zeng, K. Sycara, How can an agent learn to negotiate, in: J. Mueller, M. Wooldridge, N. Jennings (Eds.), *Intelligent Agents III, Agent Theories, Architectures and Languages, Lecture Notes in Artificial Intelligence*, vol. 1193, Springer, Berlin, 1997, pp. 233–244.



**Peyman Faratin** is a research assistant in the Intelligent Systems Group in the Department of Electronic Engineering at Queen Mary and Westfield (QMW) College (University of London). His research topic is automated negotiation mechanisms for real world applications which have included telecommunications network management (with Nortel Technology) and business process management (with British Telecom).



**Carles Sierra** is a researcher at the Institute of research on Artificial Intelligence of the Spanish Council for Scientific Research (CSIC). He received the M.S. in Computer Science (1986) and the Ph.D. in Computer Science (1989) from the Technical University of Catalonia, Barcelona, Spain. His current research interests include: multi-agent systems, formal methods, computational reflection, uncertainty, robotics, and applications of AI to medicine. He has participated in more than 20 research projects funded by the European Commission and the Spanish Government, and has published more than 80 papers in specialised conferences and scientific journals. Recently, he has been particularly active in the area of electronic commerce. He has published several papers on modelling electronic auctions, and on automatic negotiation models among autonomous agents.



**Nick R. Jennings** is head of the Intelligent Systems Group in the Department of Electronic Engineering at Queen Mary and Westfield (QMW) College (University of London). He carries out applied and basic research into a broad range of issues related to the deployment of agent technology in real-world applications. He has designed and built agent-based applications for the followings domains: telecommunications network management, business process management, electricity management, patient care, digital libraries, data interpretation and process control. He has also made a number of contributions to the study of agent interactions – including: social rationality, cooperation, coordination, negotiation and argumentation.

He has been an invited lecturer at numerous national and international conferences related to agent systems, he has initiated and co-chaired two major international conferences (The Practical Application of Agents and Multi-Agent Systems (PAAM 96) and Autonomous Agents (AA97)), and has initiated and co-chaired the Agent Theories, Architectures and Languages Workshop series. He has published some 85 articles on various facets of agent-based computing, written one monograph, and co-edited four books. He is the editor-in-chief of the *International Journal Autonomous Agents and Multi-Agent Systems*.