Multi-Issue Negotiation Under Time Constraints

Shaheen S. Fatima
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF U.K.
shaheen@csc.liv.ac.uk

Michael Wooldridge
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF U.K.
mjw@csc.liv.ac.uk

Nicholas R.Jennings
Department of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ U.K.
nrj@ecs.soton.ac.uk

ABSTRACT

This paper presents a new model for multi-issue negotiation under time constraints in an incomplete information setting. In this model the order in which issues are bargained over and agreements are reached is determined endogenously as part of the bargaining equilibrium. We show that the sequential implementation of the equilibrium agreement gives a better outcome than a simultaneous implementation when agents have like, as well as conflicting, time preferences. We also show that the equilibrium solution possesses the properties of uniqueness and symmetry, although it is not always Pareto-optimal.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; K.4.4 [**Computers and Society**]: Electronic Commerce

General Terms

Algorithms, Design

Keywords

Negotiation, game theory, agendas

1. INTRODUCTION

Agent mediated negotiation has received considerable attention in the field of electronic commerce [14, 9, 7]. In many of the applications that are conceived in this domain it is important that the agents should not only bargain over the price of a product, but also take into account aspects like the delivery time, quality, payment methods, and other product specific properties. In such multi-issue negotiations, the agents should be able to negotiate outcomes that are mutually beneficial for both parties [11]. However the complexity of the bargaining problem increases rapidly as the number of issues increases. Given this increase in complexity, there is a need to develop software agents that can operate effectively in such circumstances. To this end, this paper reports on the development of a new model for multi-issue negotiation between two agents.

In such bilateral multi-issue negotiations, one approach is to bundle all the issues and discuss them simultaneously. This allows the players to exploit trade-offs among different issues, but requires complex computations to be performed [11]. The other approach — which is computationally simpler — is to negotiate the issues sequentially. Although issue-by-issue negotiation minimizes the complexity of the negotiation procedure, an important question that arises is the order in which the issues are bargained. This ordering is called the negotiation agenda. Moreover, one of the factors that determines the outcome of negotiation is this agenda [4]. To this end, there are two ways of incorporating agendas in the negotiation model. One is to fix the agenda exogenously as part of the negotiation procedure [4]. The other way, which is more flexible, is to allow the bargainers to decide which issue they will negotiate next during the process of negotiation. This is called an endogenous agenda [5]. Against this background, this paper presents a multi-issue negotiation model with an *endogenous agenda*.

To provide a setting for our negotiation model, we consider the case in which negotiation needs to be completed by a specified time (which may be different for the different parties). Apart from the agents' respective deadlines, the time at which agreement is reached can affect the agents in different ways. An agent can gain utility with time, and have the incentive to reach a late agreement (within the bounds of its deadline). In such a case it is said to be a strong (patient) player. The other possibility is that it can lose utility with time and have the incentive to reach an early agreement. It is then said to be a weak (impatient) player. As we will show, this disposition and the actual deadline itself strongly influence the negotiation outcome. Other parameters that affect the outcome include the agents'strategies, their utilities and their reservation limits. However, in most practical cases agents do not have complete information on all of these parameters. Thus in this work we focus on bilateral negotiation between agents with *time constraints* and *incomplete information*.

To this end, Fatima et al presented a single-issue model for negotiation between two agents under time constraints and in an incomplete information setting [3]. Within this context, they determined optimal strategies for agents but did not address the issue of the existence of equilibrium. Here we adopt this framework and prove that mutual strategic behavior of agents, where both use their respective optimal strategies, results in equilibrium. We then extend this framework for multi-issue negotiation between a buyer and a seller for the price of more than one good/service. Specifically, each agent has a deadline before which agreement must end on all the issues. However, the order in which issues are bargained over and agreements are reached is determined by the equilibrium strategies. These strategies optimize the time at which an issue is settled and are therefore appropriate for the sequential implemen-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'02, July 15-19, Bologna, Italy
Copyright 2002 ACM 1-58113-480-0/02/0007 ...$5.00.

tation scheme. Moreover, we show that the sequential implementation of the equilibrium agreement results in an outcome that is no worse than the outcome for the simultaneous implementation, both when agents have like as well as conflicting time preferences. Finally, we study the properties of the equilibrium solution.

This work extends the state of the art by presenting a more realistic negotiation model that captures the following three aspects of many real life bargaining situations. Firstly, it is a model for negotiating multiple issues. Secondly, it takes the time constraints of bargainers into consideration. Thirdly it allows agents to have incomplete information about each other.

In section 2 we first give an overview of the single-issue negotiation model of [3] and then prove that the mutual strategic behavior of agents where both use their respective optimal strategies results in equilibrium. In section 3 we extend this model to allow multi-issue negotiation and study the properties of the equilibrium solution. Section 4 discusses related work. Finally, related: Command not found.

section 5 gives some conclusions.

2. SINGLE-ISSUE NEGOTIATION MODEL

In this section we first provide an overview of the single issue negotiation model and a brief description of the optimal strategies as determined in [3]. Due to lack of space, we describe the optimal strategy determination for one specific negotiation scenario. We then prove that the optimal strategy profiles form sequential equilibrium points.

2.1 The Negotiation Protocol

This is basically an alternating offers protocol. Let b denote the buyer, s the seller and let $[P^a_{min}, P^a_{max}]$ denote the range of values for price that are acceptable to agent a, where $a \in \{b, s\}$. A value for price that is acceptable to both b and s, i.e., the zone of agreement, is the interval $[P^s_{min}, P^b_{max}]$. $(P^b_{max} - P^s_{min})$ is the *price-surplus*. T^a denotes agent a's deadline. Let $p^t_{b \to s}$ denote the price offered by agent b at time t. Negotiation starts when the first offer is made by an agent. When an agent, say s, receives an offer from agent b at time t, i.e., $p^t_{b \to s}$, it rates the offer using its utility function U^s. If the value of U^s for $p^t_{b \to s}$ at time t is greater than the value of the counter-offer agent s is ready to send at time t', i.e., $p^{t'}_{s \to b}$ with $t' > t$ then agent s accepts. Otherwise a counter-offer is made. Thus the action A that agent s takes at time t is defined as:

$$A^s(t', p^t_{b \to s}) = \begin{cases} Quit & \text{if } t > T^s \\ Accept & \text{if } U^s(p^t_{b \to s}) \geq U^s(p^{t'}_{s \to b}) \\ p^{t'}_{s \to b} & \text{otherwise} \end{cases}$$

2.2 Counter-offer generation

Since both agents have a deadline, we assume that they use a time dependent tactic (e.g. linear (L), Boulware (B) or Conceder (C)) [2] for generating the offers. In these tactics, the predominant factor used to decide which value to offer next is time t. The tactics vary the value of price depending on the remaining negotiation time, modeled as the above defined constant T^a. The initial offer is a point in the interval $[P^a_{min}, P^a_{max}]$. The constant k^a multiplied by the size of interval determines the price to be offered in the first proposal by agent a (as per [2]). The offer made by agent a to agent b at time t ($0 < t \leq T^a$) is modeled as a function ϕ^a depending on time as follows:

$$p^t_{a \to b} = \begin{cases} P^a_{min} + \phi^a(t)(P^a_{max} - P^a_{min}) & \text{for } b \\ P^a_{min} + (1 - \phi^a(t))(P^a_{max} - P^a_{min}) & \text{for } s. \end{cases}$$

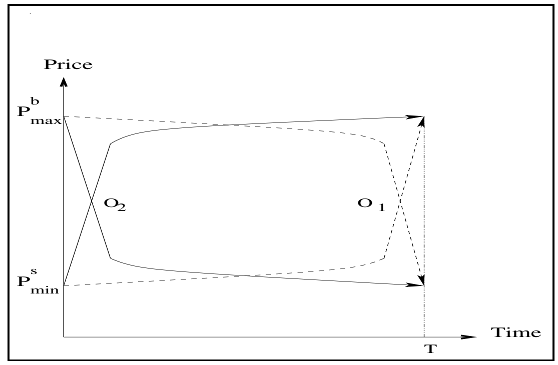

Figure 1: Negotiation outcome for Boulware and Conceder functions

A wide range of time dependent functions can be defined by varying the way in which $\phi^a(t)$ is computed (see [2] for more details). However, functions must ensure that $0 \leq \phi^a(t) \leq 1$, $\phi^a(0) = k^a$ and $\phi^a(T^a) = 1$. That is, the offer will always be between the range $[P^a_{min}, P^a_{max}]$, at the beginning it will give the initial constant and when the deadline is reached it will offer the reservation value. Function $\phi^a(t)$ is defined as follows:

$$\phi^a(t) = k^a + (1 - k^a) \left(\frac{min(t, T^a)}{T^a} \right)^{\frac{1}{\psi}}$$

These families of negotiation decision functions (NDF) represent an infinite number of possible tactics, one for each value of ψ. However, depending on the value of ψ, two extreme sets show clearly different patterns of behaviour.

1. *Boulware* [11]. For this tactic $\psi < 1$ and close to zero. The initial offer is maintained till time is almost exhausted, when the agent concedes up to its reservtion value.
2. *Conceder* [10]. For this tactic ψ is high. The agent goes to its reservation value very quickly and maintains the same offer till the deadline. Finally when $\psi = 1$ price is increased linearly.

The value of a counter offer depends on the initial price (IP) at which the agent starts negotiation, the final price (FP) beyond which the agent does not concede, ψ and T^a. A vector V of these four variables, i.e., V = [IP, FP, T^a, ψ] forms the agent's *strategy*. Let $p \in [P^b_{max}, P^s_{min}]$ and $t \in [0, T^a]$. The *negotiation outcome* O is an element of $\{(p, t), C\}$, where the pair (p, t) denotes the *price* and *time* at which agreement is reached and C denotes the conflict outcome.

For example, when b's strategy is defined as $[P^s_{min}, P^b_{max}, T, B]$ and s's strategy is defined as $[P^b_{max}, P^s_{min}, T, B]$, the outcome ($O_1$) that results is shown in Figure 1. As shown in the figure, agreement is reached at a price $P^s_{min} + (price\text{-}surplus/2)$ and at a time close to T. Similarly when the NDF in both strategies is replaced with C, then agreement (O_2) is reached at the same price but towards the beginning of negotiation.

2.3 Agents' information state

Each agent has a reservation limit, a deadline, a utility function and a strategy. Thus the buyer and seller each have four parameters denoted $\langle P^b_{max}, T^b, U^b, S^b \rangle$ and $\langle P^s_{min}, T^s, U^s, S^s \rangle$ respectively. The outcome of negotiation depends on all these eight parameters.

The information state I^a of an agent a is the information it has about the negotiation parameters. An agent's own parameters are known to it, but the information it has about the opponent is not complete. I^b and I^s are taken as:

$$I^b = \langle P^b_{max}, T^b, U^b, S^b, L^s_p, L^s_t \rangle$$

and

$$I^s = \langle P^s_{min}, T^s, U^s, S^s, L^b_p, L^b_t \rangle$$

where P^b_{max}, T^b, U^b and S^b are the information about b's own parameters and L^s_p and L^s_t are its beliefs about s. Similarly, P^s_{min}, T^s, U^s and S^s are s's own parameters and L^b_p and L^b_t are its beliefs about b. L^s_p and L^s_t are two lotteries that denote b's beliefs about s's deadline and reservation price, where
$L^s_p = [\alpha^b, P^s_{minl}; 1-\alpha^b, P^s_{minh}]$ such that $(P^s_{minl} < P^s_{minh})$ and
$L^s_t = [\gamma^b, T^s_l; 1-\gamma^b, T^s_h]$ such that $(T^s_l < T^s_h)$.
Similarly L^b_p and L^b_t are two lotteries that denote s's beliefs about b's deadline and reservation price, where
$L^b_p = [\alpha^s, P^b_{maxl}; 1-\alpha^s, P^b_{maxh}]$ such that $(P^b_{maxl} < P^b_{maxh})$ and
$L^b_t = [\gamma^s, T^b_l; 1-\gamma^s, T^b_h]$ such that $(T^b_l < T^b_h)$.
Thus agents have uncertain information about each other's deadline and reservation price. However, the agents do not know their opponent's utility function or strategy.

Agents' utilities are defined with the following two von Neumann-Morgenstern utility functions [6] that incorporate the affect of time discounting. $U^a(p,t) = U^a_p(p)U^a_t(t)$, where U^a_p and U^a_t are unidimensional utility functions. Here, preferences for attribute p, given the other attribute t, do not depend on the level of t. U^a_p is defined as:

$$U^a(p) = \begin{cases} P^b_{max} - p & \text{for the buyer} \\ p - P^s_{min} & \text{for the seller} \end{cases}$$

U^a_t is defined as $U^a_t(t) = (\delta^a)^t$, where δ^a_t is the discounting factor. Thus when $(\delta^a > 1)$ the agent is patient and gains utility with time and when $(\delta^a < 1)$ the agent is impatient and loses utility with time. Note that the agents may have different discounting factors.

Agents are said to have similar time preferences if both gain on time or both lose on time. Otherwise they have conflicting time preferences. Each agent's information is its *private* information that is not known to the opponent.

2.4 Optimal strategies

We describe how optimal strategies are obtained for players that are von Neumann-Morgenstern expected utility maximizers. Since utility is a function of price and time, these strategies optimize both. The discussion is from the perspective of the buyer (although the same analysis can be taken from the perspective of the seller). b believes that with probability γ^b, s's deadline is T^s_l and with probability $1-\gamma^b$ it is T^s_h. Let T^b denote b's own deadline. This gives rise to three relations between agent deadlines; $(T^b > T^s_h)$, $(T^s_l < T^b < T^s_h)$ or $(T^b < T^s_l)$. For each of the two possible realizations of b's discounting factor, these three relations can hold between agent deadlines. In other words, there are six possible scenarios $(N_1..N_6)$ under which negotiation can take place. Due to lack of space we describe here how the optimal strategies are obtained for one specific scenario – N_1, i.e, when b gains on time (i.e., $\delta^b > 1$) and $T^b > T^s_h$.

Since $\delta^b > 1$, b prefers to reach agreement at the latest possible time and at the lowest possible price. The optimal strategy for b is determined using *backward induction*. The optimal price (P^s_o) and the optimal time (T^s_o) are determined first, and then a strategy that ensures agreement at P^s_o and T^s_o is devised.

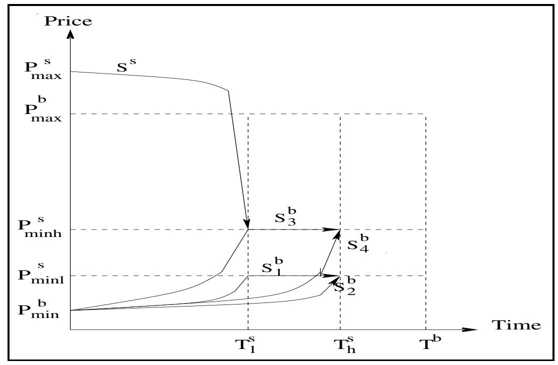

Figure 2: Possible buyer strategies in scenario N_1

No matter which strategy (B, C or L) s uses, it is bound to reach its reservation price by its deadline. Since both possible values of s's deadline are less than T^b, b's optimal strategy would be never to offer a price more than s's reservation price (P^s_{minl} or P^s_{minh}). Moreover, it should offer this price at the latest possible time, which is s's deadline (T^s_l or T^s_h). This is because s quits if agreement is not reached by then.

From its beliefs, it is known to b that s's reservation price, deadline pair could be one of (P^s_{minl}, T^s_l), (P^s_{minl}, T^s_h), (P^s_{minh}, T^s_l) or (P^s_{minh}, T^s_h). One of these four pairs is (P^s_o, T^s_o). The possible strategies that b can use are S^b_1 if s's reservation price deadline pair is (P^s_{minl}, T^s_l), S^b_2 if it is (P^s_{minl},T^s_h), S^b_3 if it is (P^s_{minh}, T^s_l) or S^b_4 if it is (P^s_{minh}, T^s_h). These strategies are depicted in Figure 2. Out of these four possible strategies, the one that results in maximum expected utility (EU) is b's optimal strategy (note that b's optimal strategy does not depend on s's strategy). The EUs from these four strategies depend on α^b and γ^b.

An agent's utility from price is independent of its utility from time, i.e, the buyer always prefers a low price to a high price, and for a given price it always prefers a late agreement to an early one. In order to simplify the process of finding the optimal strategy, we assume $\alpha^b = 1$, i.e., there is only one possible value, P^s_{minl}, for s's reservation price. This leaves us with only two strategies, S^b_1 and S^b_2 (see Figure 2). The EUs[1] from these strategies are:

$$EU_1 = \gamma^b U^b(P^s_{minl}, T^s_l) + (1-\gamma^b)U^b(P^s_{minl}, t)$$

where $T^s_l < t \leq T^s_h$ and

$$EU_2 = \gamma^b U^b(Conflict) + (1-\gamma^b)U^b(P^s_{minl}, T^s_h)$$

Out of these two, the one that gives a higher utility is optimal. EU_1 and EU_2 depend on the value of γ^b. So γ^b is varied between 0 and 1 and EU_1 and EU_2 are computed for different values of γ^b. A comparison of these two utilities shows that for a particular value of γ^b, say γ^b_c, $EU_1 = EU_2$. For $(\gamma^b < \gamma^b_c)$, $EU_2 > EU_1$ and for $(\gamma^b > \gamma^b_c)$, $EU_1 > EU_2$. Thus γ_c is crucial in determining the optimal strategy. This computation gives the optimal time (T^s_o) for reaching agreement. T^s_o is T^s_l if $(\gamma^b > \gamma^b_c)$ or T^s_h if $(\gamma^b < \gamma^b_c)$. The next step is to find the optimal price. Assume that $(\gamma^b > \gamma^b_c)$ which means that S^b_1 is better than S^b_2. This implies that the optimal time for reaching agreement is T^s_l and not T^s_h. Strategies S^b_2 and S^b_4 can result in an agreement only after T^s_l, since the price offered prior to T^s_l is unacceptable to the seller. Thus neither S^b_2 nor S^b_4 can be optimal. The optimal strategy is S^b_1 or S^b_3 depending on the

[1] Utility from conflict to both agents is less than zero.

Negotiation Scenario	Optimal Strategy
N_1	$[P^b_{min}, P^s_o, T^s_o, B]$ for all values of t
N_2	$[P^b_{min}, P^s_o, T^s_l, B]$ if $t \leq T^s_l$ $[P^s_o, P^b_{max}, (T^b - T^s_l), B]$ if $t > T^s_l$
N_3	$[P^b_{min}, P^b_{max}, T^b, B]$ for all values of t
N_4	$[P^b_{min}, P^s_o, T^s_l, C]$ for all values of t
N_5	$[P^b_{min}, P^s_o, T', C]$ if $t \leq T^s_l$ $[P^s_o, P^b_{max}, (T^b - T^s_l), C]$ if $t > T^s_l$
N_6	$[P^b_{min}, P^b_{max}, T^b, C]$ for all values of t

Table 1: Optimal strategies for the buyer. T' denotes the beginning of negotiation.

value of α^b. The expressions for EU_1 and EU_3 are:

$$EU_1 = \alpha^b U^b(P^s_{minl}, T^s_l) + (1 - \alpha^b) U^b(Conflict)$$

and

$$EU_3 = \alpha^b U^b(p, t) + (1 - \alpha^b) U^b(P^s_{minh}, T^s_l)$$

where $(P^s_{minl} < p < P^s_{minh})$ and $(T < t < T^s_l)$. Here T denotes the time at which b offers P^s_{minl}. α^b is varied between 0 and 1 and EU_1 and EU_3 are computed. For $(\alpha^b = \alpha^b_c)$ $EU_1 = EU_3$, for $(\alpha^b < \alpha^b_c)$ $EU_3 > EU_1$ and for $(\alpha^b > \alpha^b_c) EU_1 > EU_3$. This gives the optimal price P^s_o for reaching agreement. P^s_o is P^s_{minl} if $(\alpha^b > \alpha^b_c)$ or P^s_{minh} if $(\alpha^b < \alpha^b_c)$. Assume that $(\alpha^b < \alpha^b_c)$. This means that the optimal price is P^s_{minh} and the optimal strategy is $S^b_3 = [P^b_{min}, P^s_{minh}, T^s_l, B]$. Thus S^b_3 results in an outcome that is optimal in both the price (P^s_{minh}) and time (T^s_l).

Assume that the seller also gains on time and $T^s = T^s_l$ and $P^s_{min} = P^s_{minh}$. Let the values of T^b_l and T^b_h in s's lottery be T^b and some value greater than T^b respectively. Since both possible values of the buyer's deadline are greater than its own, irrespective of its value for γ^s, it has to concede up to P^s_{min} by T^s. Thus from the seller's perspective, the optimal price $P^b_o = P^s_{min}$ and the optimal time $T^b_o = T^s$. In such a scenario, the optimal strategy for s is to start at some high price, make small concessions till its deadline is almost reached and then offer the reservation price P^s_{min} at T^s using the Boulware NDF, i.e., $S^s = [P^s_{max}, P^s_{min}, T^s, B]$. In order for the b and s strategies to converge, the values of α^b and γ^b in b's lotteries should be such that $(\alpha^b < \alpha^b_c)$ and $(\gamma^b > \gamma^b_c)$. When these conditions are satisfied P^s_o equals P^b_o and T^s_o equals T^b_o. The optimal strategies $S^b_3 \times S^s$ then converge and result in an agreement at price P^s_{minh} and time T^s_l (see Figure 2). b gets the entire price-surplus.

When both buyer and seller lose utility on time, the optimal strategy for them is to offer P^s_{minh} at the earliest opportunity. This can be done using a Conceder NDF that results in agreement at the same price P^s_{minh} but towards the beginning of negotiation. In the same way, optimal strategies are obtained for the remaining negotiation scenarios. These are summarised in Table 1. A similar kind of analysis is made from the seller's perspective to obtain P^b_o and T^b_o in the six possible scenarios. In each of these scenarios, the agents' optimal strategies do not depend on their opponent's strategy. Again see [3] for details.

There are many scenarios in which negotiation can take place. These depend on the agents' attitude towards time and the relationship between their deadlines. As stated earlier in this section, there are six possible scenarios from the buyer's perspective, on the basis of which it selects its strategy. Similarly from the seller's perspective there are also six possible scenarios. But the negotiation

Deadline Ordering	Seller's Deadline	Case 1, 2 and 3 Outcome		Case 4 Outcome	
		b's deadline		b's deadline	
D1	T^s_l	(P^s_o, T^s_l)	(P^s_o, T^s_l)	(P^s_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
	T^s_h	(P^s_o, T^s_h)	(P^s_o, T^s_h)	(P^s_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
D2	T^s_l	(P^b_o, T^b_l)	(P^b_o, T^b_h)	(P^b_o, T')	(P^b_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
	T^s_h	(P^b_o, T^b_l)	(P^b_o, T^b_h)	(P^b_o, T')	(P^b_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
D3	T^s_l	(P^s_o, T^s_l)	(P^s_o, T^s_l)	(P^s_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
	T^s_h	(P^b_o, T^b_l)	(P^s_o, T^s_h)	(P^b_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
D4	T^s_l	(P^s_o, T^s_l)	(P^s_o, T^s_l)	(P^s_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
	T^s_h	(P^b_o, T^b_l)	(P^b_o, T^b_h)	(P^b_o, T')	(P^b_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
D5	T^s_l	(P^b_o, T^b_l)	(P^s_o, T^s_l)	(P^b_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
	T^s_h	(P^b_o, T^b_l)	(P^s_o, T^s_h)	(P^b_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
D6	T^s_l	(P^b_o, T^b_l)	(P^s_o, T^s_l)	(P^b_o, T')	(P^s_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h
	T^s_h	(P^b_o, T^b_l)	(P^b_o, T^b_h)	(P^b_o, T')	(P^b_o, T')
		T^b_l	T^b_h	T^b_l	T^b_h

Table 2: Outcome of negotiation when both agents use their respective optimal strategies

outcome depends on all possible ways in which interaction between b and s can take place. There can be six possible orderings on the agent deadlines:

1. $T^s_l < T^s_h < T^b_l < T^b_h$ (D1)
2. $T^b_l < T^b_h < T^s_l < T^s_h$ (D2)
3. $T^s_l < T^b_l < T^s_h < T^b_h$ (D3)
4. $T^s_l < T^b_l < T^b_h < T^s_h$ (D4)
5. $T^b_l < T^s_l < T^s_h < T^b_h$ (D5)
6. $T^b_l < T^s_l < T^b_h < T^s_h$ (D6)

For each of these orderings, the agents' attitudes towards time could be one of the following:

1. Both buyer and seller gain utility with time (Case 1).
2. Buyer gains and seller loses utility with time (Case 2).
3. Buyer loses and seller gains utility with time (Case 3).
4. Both buyer and seller lose utility with time (Case 4).

Thus in total there are 24 possible negotiation scenarios and the outcome of negotiation depends on the exact scenario. A summary of these is given in Table 2. P^s_o indicates that the price-surplus goes to b and P^b_o indicates that the price-surplus goes to s. As seen in this table, the price-surplus always goes to the agent with the longer deadline. The time of agreement is T' (which denotes

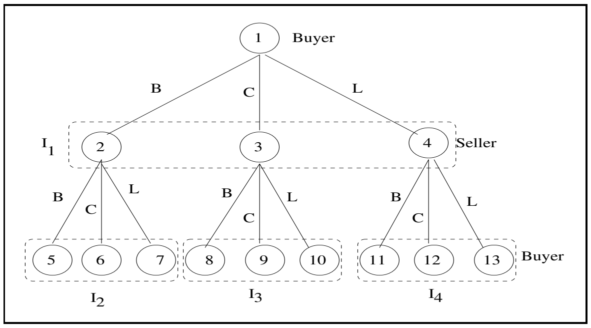

Figure 3: Extensive form of the negotiation game

the beginning of negotiation) if both agents lose on time, and the earlier deadline if at least one agent gains on time. Note that these are the outcomes that will result if the agents' beliefs about each other satisfy the following conditions for convergence of strategies.

1. $(\gamma^b < \gamma_c^b)$ if $(T^s = T_h^s)$ and $(\gamma^b > \gamma_c^b)$ if $(T^s = T_l^s)$ for b.
2. $(\alpha^b < \alpha_c^b)$ if $(P_{min}^s = P_{minh}^s)$ and $(\alpha^b > \alpha_c^b)$ if $(P_{min}^s = P_{minl}^s)$ for b.
3. $(\gamma^s < \gamma_c^s)$ if $(T^b = T_h^b)$ and $(\gamma^s > \gamma_c^s)$ if $(T^b = T_l^b)$ for s.
4. $(\alpha^s < \alpha_c^s)$ if $(P_{max}^b = P_{maxh}^b)$ and $(\alpha^s > \alpha_c^s)$ if $(P_{max}^b = P_{maxl}^b)$ for s.

The similarity between these results and those of Sandholm and Vulkan [15] on bargaining with deadlines is that, in both cases, the price-surplus always goes to the agent with the longer deadline. However, the difference is that in [15] the deadline affect overrides time discounting, whereas here the deadline affect does not override time discounting. This happens because in [15] the agents always make offers that lie within the zone of agreeement. In our model, agents initially make offers that lie outside this zone, and thereby delay the time of agreement. Thus when agents have conflicting time preferences, in our case, agreement is reached near the earlier deadline, but in [15] agreement is reached towards the beginning of negotiation.

The single issue negotiation model of [3] only determines optimal strategies for agents on the basis of available information and shows the resulting outcome. However such an outcome is only possible if this mutual strategic behavior of agents leads to equilibrium. In the following subsection we prove this by using the standard game theoretic solution concept of sequential equilibrium.

2.5 Equilibrium agreements

Since agents do not have information about their opponent's strategy or utility, negotiation can be considered as a game $\boldsymbol{G}$ of incomplete information. A *strategy profile* and *belief system* pair is a *sequential equilibrium* of an extensive game if it is *sequential rational* and *consistent* [8]. A system of beliefs μ in an extensive form game $\boldsymbol{G}$ is a specification of a probability $x \in [0, 1]$ for each decision node x in $\boldsymbol{G}$ such that $\sum_{x \in \boldsymbol{I}} \mu(x) = 1$ for all information sets $\boldsymbol{I}$. In other words, μ represents the agent's beliefs about the *history* of negotiation. The player's strategies satisfy sequential rationality if for each information set of each player a, a's strategy is a best response to the other player's strategies, given a's beliefs at that information set. The requirement for μ to be consistent with the strategy profile is as follows. Even at an infromation set that is not reached if all players adhere to their strategies, it is required that a player's belief be derived from some strategy profile using Bayes' rule.

THEOREM 1. *There exists sequential equilibrium of $\boldsymbol{G}$ at the point $[P_{min}^b, P_o^s, T_o^s, B] \times [P_{max}^s, P_o^b, T_o^b, B]$ for the negotiation scenario corresponding to case 1 and deadline ordering $D1$, where $P_o^s = P_o^b = P_{minl}^s$ if $(P_{min}^s = P_{minl}^s)$ or $P_o^s = P_o^b = P_{minh}^s$ if $(P_{min}^s = P_{minh}^s)$ and $T_o^s = T_o^b = T_l^s$ if $(T^s = T_l^s)$ or $T_o^s = T_o^b = T_h^s$ if $(T^s = T_h^s)$.*

PROOF. *The first three levels of the extensive form of this game are shown in Figure 3. At node 1 one of the players, say b, starts negotiation by using its optimal strategy $[P_{min}^b, P_o^s, T_o^s, B]$. Play reaches node 2. At this level it is player s's turn to make a decision. $\boldsymbol{I}_1$ becomes the information set for s since it is unaware of the strategy used by b and hence does not know which of the three nodes 2, 3 or 4 play has reached. However, irrespective of which node play reaches at this level (i.e., irrespective of s's belief about the history of negotiation), the dominant strategy for s is $[P_{max}^s, P_o^b, T_o^b, B]$. Play now reaches node 5 (since both agents use B) at which b makes a move. At this point b does not know exactly which node the play is at, but it knows with certainty that its information set $\boldsymbol{I}_2$ is reached (the probability of reaching other decision nodes at this level is 0). The dominant strategy for b at this information set (and at all others) is $[P_{min}^b, P_o^s, T_o^s, B]$. Thus at every information set at which it is b's turn to make a move, its optimal strategy is $[P_{min}^b, P_o^s, T_o^s, B]$, and at every information set at which it is s's turn to make a move, its optimal strategy is $[P_{max}^s, P_o^b, T_o^b, B]$. The strategy profile $[P_{min}^b, P_o^s, T_o^s, B] \times [P_{max}^s, P_o^b, T_o^b, B]$ therefore satisfies the requirements for sequential rationality. Furthermore, at every information set, the optimal strategies are also dominant strategies. This makes the strategy profile $[P_{min}^b, P_o^s, T_o^s, B] \times [P_{max}^s, P_o^b, T_o^b, B]$ a sequential equilibrium point irrespective of the agents' beliefs about the history of negotiation.* □

COROLLARY 1. *The optimal strategy profile constitutes a unique equilibrium.*

PROOF. *This is a direct consequence of the above proof. As the optimal strategies for both agents are dominant strategies at each of their information sets, there does not exist any other equilibrium (neither a pure nor a mixed strategy) where an agent uses a strategy other than its optimal strategy.* □

In the same way, sequential equilibrium can be shown to exist when agents use their optimal strategies in all the remaining negotiation scenarios.

3. MULTI-ISSUE NEGOTIATION

We now extend the above model for multi-issue bargaining where the issues are independent[2] of each other. Assume that buyer, b, and seller, s, that have unequal deadlines, bargain over the price of two distinct goods/services, X and Y. Negotiation on all the issues must end before the deadline. We consider two goods/services in order to simplify the discussion but this is a general framework that works for more than two goods/services.

3.1 Agents' information state

Let the buyer's reservation prices for X and Y be P_x^b and P_y^b and the seller's reservation prices be P_x^s and P_y^s respectively. The buyer's information state is:

$$I^b = \langle P_x^b, P_y^b, T^b, U^b, S^b, L_x^s, L_y^s, L_t^s \rangle$$

[2]Independence is a common and reasonable assumption to make in this context. Future work will deal with the dependent case.

where P_x^b, P_y^b, T^b, U^b and S^b are the information about its own parameters and L_x^s, L_y^s and L_t^s are three lotteries that denote its beliefs about the opponent's parameters.
$L_x^s = [\alpha^b, P_{xL}^s; 1-\alpha^b, P_{xH}^s]$ is the lottery on the seller's reservation price for X such that $P_{xL}^s < P_{xH}^s$,
$L_y^s = [\beta^b, P_{yL}^s; 1-\beta^b, P_{yH}^s]$ is the lottery on the seller's reservation price for Y such that $P_{yL}^s < P_{yH}^s$ and
$L_t^s = [\gamma^b, T_l^s; 1-\gamma^b, T_h^s]$ is the lottery on the seller's deadline such that $T_l^s < T_h^s$.
Similarly, the seller's information state is defined as:

$$I^s = \langle P_x^s, P_y^s, T^s, U^s, S^s, L_x^b, L_y^b, L_t^b \rangle$$

An agent's information state is its private knowledge. The agents' utility functions are defined as:

$$U^a(p_x, p_y, t) = \begin{cases} (P_x^b - p_x)(\delta_x^b)^t + (P_y^b - p_y)(\delta_y^b)^t & \text{for } b \\ (p_x - P_x^s)(\delta_x^s)^t + (p_y - P_y^s)(\delta_y^s)^t & \text{for } s \end{cases}$$

Note that the discounting factors are different for different issues. This allows agents' attitudes toward time to be different for different issues.

3.2 Negotiation protocol

Again we use an alternating offers negotiation protocol. There are two types of offers. An offer on just one good is referred to as a *single offer* and an offer on two goods is referred to as a *combined offer*. One of the agents starts by making a combined offer. The other agent can accept/reject part of the offer (single issue) or the complete offer. If it rejects the complete offer, then it sends a combined counter-offer. This process of making combined offers continues till agreement is reached on one of the issues. Thereafter agents make offers only on the remaining issue (i.e., once agreement is reached on an issue, it cannot be renegotiated). Negotiation ends when agreement is reached on both the issues or a deadline is reached. Thus the action A that agent s takes at time t on a single offer is as defined in section 2.1 . Its action on a combined offer, $A^s(t', X_{b\to s}^t, Y_{b\to s}^t)$, is defined as:

1. Quit if $t > T^s$
2. Accept $X_{b\to s}^t$ if $U^s(X_{b\to s}^t, t) \geq U^s(X_{s\to b}^{t'}, t')$
3. Accept $Y_{b\to s}^t$ if $U^s(Y_{b\to s}^t, t) \geq U^s(Y_{s\to b}^{t'}, t')$
4. Offer $X_{s\to b}^{t'}$ if $X_{b\to s}^t$ not accepted
5. Offer $Y_{s\to b}^{t'}$ if $Y_{b\to s}^t$ not accepted.

A counter-offer for an issue is generated using the method described in section 2.2. Although agents initially make offers on both issues, there is no restriction on the price they offer. Thus by initially offering a price that lies outside the zone of agreement, an agent can effectively delay the time of agreement for that issue. For example, b can offer a very low price which will not be acceptable to s and s can offer a price which will not be acceptable to b. In this way, the order in which the issues are bargained over and agreements are reached is determined endogenously as part of the bargaining equilibrium rather than imposed exogenously as part of the game tree.

Two implementation rules are possible for this protocol [4]. One is *sequential implementation* in which agreement on an issue is implemented as soon as it is settled; and the other is *simultaneous implementation* in which agreement is implemented only after all the issues are settled. We first list the equilibrium agreements in different negotiation scenarios and then compare the outcome that results from the sequential implementation with that of the simultaneous implementation.

	Negotiation Scenario		Time of agreement	
	b	s		
	XY	XY	X	Y
1	GG	GG	T	T
2	GG	GL	T	T
3	GG	LG	T	T
4	GG	LL	T	T
5	GL	GG	T	T
6	GL	GL	T	T'
7	GL	LG	T	T
8	GL	LL	T	T'
9	LG	GG	T	T
10	LG	GL	T	T
11	LG	LG	T'	T
12	LG	LL	T'	T
13	LL	GG	T	T
14	LL	GL	T	T'
15	LL	LG	T'	T
16	LL	LL	T'	T'

Table 3: Equilibrium agreements for two issues X and Y

3.3 Equilibrium strategies

We assume that the conditions for convergence (as listed in section 2.4) are satisfied for both X and Y. As agents negotiate over the price of two distinct goods/services, the equilibrium strategies for the single issue model can be applied to X and Y independently of each other. The equilibrium agreements in different negotiation scenarios are listed in Table 3. T equals T^b if ($T^b < T^s$) and T^s if ($T^s < T^b$). T' denotes the beginning of negotiation. G indicates that the agent gains utility on time and L indicates that it loses on time. The price-surplus on both issues always goes to the agent with the longer deadline (see section 2.4).

Consider a situation where both b and s lose on time on X and gain on time on Y (row 11 of Table 3). Let $T^b > T^s$. Assuming the conditions for convergence are satisfied, b's equilibrium strategies for X and Y are
$[P_{xmin}^b, P_x^s, T^s, C]$ and $[P_{ymin}^b, P_y^s, T^s, B]$
and those for s are
$[P_{xmax}^s, P_x^s, T^s, C]$ and $[P_{ymax}^s, P_y^s, T^s, B]$
(see Table 1). During the process of negotiation, agents generate offers using these strategies. This results in an agreement on X towards the beginning of negotiation, and on Y at time T^s (which is the earlier deadline). The price-surplus for X and Y goes to the agent with the longer deadline, i.e., b.

3.4 Implementation schemes

Any two strategies (S^b, S^s) lead to an outcome of the game. If S^b and S^s are the equilibrium strategies, then the outcome is an agreement on X at time t and price p_x and an agreement on Y at time τ and price p_y. Payoffs for this outcome depend on the rules by which agreements are implemented.

- Sequential implementation. Exchange of a given good/service takes place at the time of agreement on a price for that good/service. The agents' utilities from the strategy pair (S^b, S^s) leading to agreements (p_x, t) and (p_y, τ) are:
$U_{seq}^b((p_x, t), (p_y, \tau)) = (P_x^b - p_x)(\delta_x^b)^t + (P_y^b - p_y)(\delta_y^b)^\tau$
$U_{seq}^s((p_x, t), (p_y, \tau)) = (p_x - P_x^s)(\delta_x^s)^t + (p_y - P_y^s)(\delta_y^s)^\tau$
- Simultaneous implementation. Exchange of goods/services takes place only after agreement is reached on prices of all

the goods. The agents'utilities for this rule are:
$U^b_{sim}((p_x, t), (p_y, \tau)) = (P^b_x - p_x)(\delta^b_x)^{max(t,\tau)} + (P^b_y - p_y)(\delta^b_y)^{max(t,\tau)}$
$U^s_{sim}((p_x, t), (p_y, \tau)) = (p_x - P^s_x)(\delta^s_x)^{max(t,\tau)} + (p_y - P^s_y)(\delta^s_y)^{max(t,\tau)}$

Since the equilibrium strategies optimize the time (and price) of agreement on an issue, it seems obvious that the agents will be better off if exchange takes place sequentially rather than simultaneously. However, since agents can have like, as well as, conflicting time preferences, it is important to determine if sequential implementation proves better than simultaneous implementation for both agents under all negotiation scenarios. We show below that sequential implementation of the equilibrium agreement always gives a better outcome than simultaneous implementation.

THEOREM 2. *The outcome generated by sequential implementation is no worse than the outcome for simultaneous implementation for both agents.*

PROOF. *When at least one of the agents gains on time on an issue, say X, then the equilibrium strategies result in an agreement at the earlier of the two deadlines. If $T^b > T^s$, then $t = T^s$ and $p_x = P^s_x$ and if $T^s > T^b$ then $t = T^b$ and $p_x = P^b_x$. When both agents lose on time on X, then agreement is reached towards the beginning of negotiation. Thus $t = T'$ and $p_x = P^s_x$ if $T^b > T^s$ and $p_x = P^b_x$ if $T^s > T^b$. As shown in Table 3, the agents have like time preferences in the first and last rows. In all other cases they have conflicting preferences on at least one issue. There are three possible ways in which agreement can be reached between agents. We analyze each of these cases below.*

1. *Both issues are agreed upon near the earlier deadline. Here $t = \tau$. Such an agreement is possible only if, for every issue, at least one agent gains on time. If p_x and p_y are the prices that are agreed for X and Y, the expressions below yield equal utility from both implementation schemes to both b and s.*

$$U^b_{seq} = U^b_{sim} = (P^b_x - p_x)(\delta^b_x)^\tau + (P^b_y - p_y)(\delta^b_y)^\tau$$

$$U^s_{seq} = U^s_{sim} = (p_x - P^s_x)(\delta^s_x)^\tau + (p_y - P^s_y)(\delta^s_y)^\tau$$

2. *Both issues are agreed upon towards the beginning of negotiation. This happens when both agents lose on time on both the issues. As in case 1, the expressions for utility from sequential and simultaneous schemes yield the same values since $t = \tau = T'$.*

3. *One issue is agreed towards the beginning of negotiation and the other near the earlier deadline. This occurs when both agents lose on time on one of the issues, say X, and at least one agent gains on time on the other issue, say Y. Here $t = T'$ and the buyer's utilities from sequential and simultaneous implementations are:*

$$U^b_{seq} = (P^b_x - p_x)(\delta^b_x)^{T'} + (P^b_y - p_y)(\delta^b_y)^\tau$$

$$U^b_{sim} = (P^b_x - p_x)(\delta^b_x)^\tau + (P^b_y - p_y)(\delta^b_y)^\tau.$$

The utility from X is greater for sequential implementation since $T' < \tau$ and both agents lose on time. The utility from Y is equal for both schemes. As a result, sequential implementation gives a total utility that is higher than simultaneous implementation. The utility that the seller gets is:

$$U^s_{seq} = (p_x - P^s_x)(\delta^s_x)^{T'} + (p_y - P^s_y)(\delta^s_y)^\tau$$

$$U^s_{sim} = (p_x - P^s_x)(\delta^s_b)^\tau + (p_y - P^s_y)(\delta^s_y)^\tau.$$

As s also loses on time on X, its utility from X is higher for sequential implementation giving a higher cumulative utility than simultaneous implementation.

Thus sequential implementation always gives a better outcome than simultaneous implementation. □

The same argument holds good when b and s negotiate over more than two issues. Thus from the perspective of both agents, sequential implementation proves to be a better implementation scheme than simultaneous implementation.

3.5 Properties of the equilibrium solution

The main focus in the design of a negotiation model is on the properties of the outcome, since the choice of a model depends on the attributes of the solution it generates. We therefore study some important properties [8] of the equilibrium agreement.

1. *Uniqueness*. If the solution of the negotiation game is unique, then it can be identified unequivocally.

   THEOREM 3. *For each negotiation scenario, the proposed negotiation model has a unique equilibrium agreement.*

   PROOF. *There are n independent negotiation issues each of which has a single equilibrium agreement (see section 2.5 for proof). This gives a unique equilibrium agreement for all n issues.* □

2. *Symmetry*. A bargaining mechanism is said to be symmetric if it does not treat the players differently on the basis of inappropriate criteria. Exactly what constitutes inappropriate criteria depends on the specific domain. The proposed negotiation mechanism possesses the property of symmetry since the outcome does not depend on which player starts the process of negotiation.

   THEOREM 4. *In all negotiation scenarios, the bargaining outcome is independent of the identity of the first player.*

   PROOF. *As shown in Table 3, there are two time points at which agreements can be reached; T' which denotes the beginning of negotiation or T which is the earlier deadline. At these time points one of the agents (either b or s depending on whose turn it is) offers the equilibrium solution which the other agent accepts.* □

3. *Efficiency*. An agreement is efficient if there is no wasted utility, i.e, the agreement satisfies Pareto-optimality. The equilibrium solution in the proposed model is Pareto-optimal in some, but not all, negotiation scenarios.

   THEOREM 5. *When players have opposite time preferences on an issue and the agent with the longer deadline loses on time on that issue, the equilibrium solution is not necessarily Pareto-optimal.*

   PROOF. *Consider row 5 of Table 3. Assume that $T^s < T^b$. On issue Y, the agents have conflicting time preferences. As $T^s < T^b$, $\tau = T^s$ and the price-surplus in the equilibrium solution goes to b ($p_y = P^s_y$). The utility to an agent can be increased by changing price or time or both. Price (p_y) can only be increased and τ can only be decreased, since a decrease in price or an increase in τ will be unacceptable to s. An increase in p_y decreases U^b and increases*

U^s. A decrease in τ increases U^b and decreases U^s. But a change in both p_y and τ can improve both U^b and U^s. The same argument holds for the other cases. □

In all the remanining scenarios it can be seen that the solution is Pareto-efficient; an increase in one agent's uitlity lowers its opponent's utility.

4. RELATED WORK

Fershtman [4] extends Rubinstein's complete information model [12] for splitting a single pie to multiple pies. This model imposes an agenda exogenously, and studies the relation between the agenda and the outcome of the bargaining game. It is based on the assumption that both players have identical discounting factors and does not consider agent deadlines. Similar work in a complete information setting includes [5] but it makes the agenda endogenous.

Bac and Raff [1] developed a model that has an endogenous agenda. They extend Rubinstein's model [13] for single pie bargaining with incomplete information by adding a second pie. In this model the price-surplus is known to both agents. For both agents, the discounting factor is assumed to be equal over all the issues. One of the players knows its own discounting factor and that of its opponent. The other player knows its own discounting factor but is uncertain of the opponent's discounting factor. This can take one of two values, δ_H with probability Π and δ_L with probability $1 - \Pi$. These probabilities are *common knowledge*. Thus agents have asymmetric information about discounting factors. However they do not associate deadlines with players.

The difference between these models and ours is that firstly, our model considers both agent deadlines and discounting factors. Secondly, in our case the players are uncertain about the opponent's reservation price and deadline. Each agent knows its own reservation price and deadline but has a binary probability distribution over its opponent's reservation price and deadline. Moreover, the discounting factor is different for different issues and the players have no information about the opponent's discounting factors. Thirdly, each agent's information state is its *private knowledge* which is not known to its opponent. Our model is therefore closer to most real life bargaining situations than the other models. The fourth point of difference lies in the attributes of the solution. Comparing the solution properties of these models, we see that the existing models do not have a unique equilibrium solution. The equilibrium solution depends on the identity of the first player. In our model, the equilibrium solution is unique and is independent of the identity of the first player. However, as is the case with our model, the equilibrium solution is not always Pareto-optimal in the other models.

5. CONCLUSIONS

This paper presented a model for multi-issue negotiation under time constraints in an incomplete information setting. The order in which issues are bargained over and agreements are reached is determined endogenously, as part of the bargaining equilibrium rather than imposed exogenously, as part of the game tree. An important property of this model is the existence of a unique equilibrium. For any issue, this equilibrium results in agreement at the earlier deadline if at least one agent has the incentive to reach a late agreement and at the beginning of negotiation if both agents have the incentive to reach an early agreement. The price-surplus on all issues goes to the agent with the longer deadline.

The sequential implementation of the equilibrium agreement was shown to result in an outcome that is no worse than the outcome for simultaneous implementation when agents have similar, as well as conflicting, time preferences. The equilibrium agreement possesses the properties of being unique and symmetric, although it is not always Pareto-optimal.

As it currently stands, our model considered the negotiation issues to be independent of each other. In future we intend to study bargaining over interdependent issues. Apart from this, our model considered the case where agents had uncertain information about each other's deadline and reservation price. In future we will introduce learning into the model that will allow the agents to learn these parameters during negotiation. These extensions will take the model further towards real life bargaining situations.

Acknowledgements

This research was supported by the EPSRC under grant GR/M07052.

6. REFERENCES

[1] M. Bac and H. Raff. Issue-by-issue negotiations: the role of information and time preference. *Games and Economic Behavior*, 13:125–134, 1996.

[2] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.

[3] S. S. Fatima, M. Wooldridge, and N. R. Jennings. Optimal negotiation strategies for agents with incomplete information. In *ATAL-2001*, pages 53–68, Seattle, USA, 2001.

[4] C. Fershtman. The importance of the agenda in bargaining. *Games and Economic Behavior*, 2:224–238, 1990.

[5] R. Inderst. Multi-issue bargaining with endogenous agenda. *Games and Economic Behavior*, 30:64–82, 2000.

[6] R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley, 1976.

[7] A. Lomuscio, M. Wooldridge, and N. R. Jennings. A classification scheme for negotiation in electronic commerce. In F. Dignum and C. Sierra, editors, *Agent-Mediated Electronic Commerce: A European Agentlink Perspective.*, pages 19–33. Springer Verlag, 2001.

[8] M.J.Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, England, 1998.

[9] R. P. Maes and A.G.Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.

[10] D. G. Pruitt. *Negotiation Behavior*. Academic Press, 1981.

[11] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, USA, 1982.

[12] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, January 1982.

[13] A. Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica*, 53:1151–1172, January 1985.

[14] T. Sandholm. Agents in electronic commerce: component technologies for automated negotiation and coalition formation. *Autonomous Agents and Multi-Agent Systems*, 3(1):73–96, 2000.

[15] T. Sandholm and N. Vulkan. Bargianing with deadlines. In *AAAI-99*, pages 44–51, Orlando, FL, 1999.