Optimal Agendas for Multi-Issue Negotiation

Shaheen 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

There are two ways of handling bilateral multi-issue negotiations – one is to negotiate all the issues together, and the other is to negotiate them one by one. The order in which issues are negotiated in issue-by-issue negotiation is specified by the agenda, which can be defined in two ways. One way is to decide it *exogenously*, i.e., before negotiation begins. The other way is to let the players decide which issue they will negotiate next, during the process of negotiation, i.e., the agenda is determined *endogenously*. Against this background, this paper studies the effect of combining the exogenous and endogenous agendas on the players' utilities. More specifically, we determine whether, decomposing a set of N issues into k stages (for $1 \leq k \leq N$), determining the issues to be negotiated at each stage exogenously, and negotiating each stage sequentially using an endogenous agenda can improve an agent's utility relative to the utility it gets if the agenda for all the N issues is defined endogenously. For each agent, we find the expected utility for each value of k between 1 and N. The value of k that gives an agent maximum utility is its optimal number of stages.

Our study shows that, in some negotiation scenarios, the optimal value of k is identical for the two players, and is greater than one. In other words, in some negotiation scenarios, both the agents can improve their utilities by using the k-stage negotiation relative to the single stage negotiation. However, since the players have incomplete information about the negotiation parameters, they cannot identify such scenarios. We therefore present an extended alternating offers protocol, that allows the agents to identify such scenarios through a *mediator*, thereby resulting in improved utility to both the agents.

Categories and Subject Descriptors

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

General Terms

Algorithms, Design

AAMAS'03, July 14-18, 2003, Melbourne, Australia

Keywords

Multi-issue negotiation, Game-theory, Agendas.

1. INTRODUCTION

In many bargaining situations the parties involved negotiate more than one issue. A concrete example of such multi-issue bargaining is the purchase of a good/service where the agents have to agree, among other things, on the price, the quality of service and financing. There are two approaches to multi-issue negotiation [8]. One is to bundle all the issues and discuss them simultaneously. This complete package approach allows the players to exploit trade-offs among different issues, but requires complex computations to be performed [10]. The other approach, which is computationally simpler, is to negotiate the issues one by one [3]. A second and perhaps more important reason why parties may choose to settle issues one by one is the strategic implications of the choice of the negotiation procedure (i.e., issue-by-issue vs complete package). When there are two issues to negotiate, the decision to negotiate them simultaneously, or one by one, is by no means neutral to the outcome [17]. 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* [8].

It has been shown in [8] that the agenda is one of the factors that determines the outcome of negotiation. For instance, if there are two issues, X and Y, the two agendas XY and YX can lead to two different outcomes. The agents need not have identical preferences over these outcomes, and one of them may prefer the agenda XY to YX, while the other may prefer YX to XY. This is because the bargainers may have time constraints for reaching agreements on the issues. Each player may have its own deadline for reaching agreement on these issues. Moreover, the endpoint may not be the only way in which time influences negotiation behaviour [14, 15, 16, 7]. Consider the case in which the service is provided immediately after the successful negotiation on an issue ends (say at price P and time T). In some situations, it is not sufficient merely for an agent to ensure that T is earlier than its deadline. This may be the case, for instance, because one of the agents, say the buyer, could be losing utility with time as a result of not getting the service. On the other hand, the seller may perhaps gain more utility by providing the service as late as possible. Thus, in this case, the seller tries to maximize T (within the limit of its deadline) and the buyer tries to minimize T. In short, agents can have different attitudes towards time. Generally speaking, the most common time effects in bargaining situatons are *time discounting* and *deadlines* [12, 11]. An agent that gains utility with time and has the incentive to reach a late agreement (within its deadline) is said to be a *patient* player.

An agent that loses utility with time and tries to reach an early agreement is said to be an *impatient* player. As we will show, this disposition and the actual deadline itself strongly influence strategic behaviour and, consequently, the negotiation outcome. Given this fact, exploring the role of the bargaining agenda, and how players might manipulate it, becomes an interesting and important area of research, especially given that many real life negotiations involve multiple issues.

In this paper we focus on procedures that allow parties to negotiate the issues successively rather than as one bundle. Issue-by-issue negotiation has been explored in [5, 6, 8, 9]. In these models the negotiation agenda is defined either *exogenously* or *endogenously*. If the agenda is defined before negotiation begins, then it is said to be exogenous. On the other hand, if the agents are allowed to decide which issue they will negotiate next, during the process of negotiation, then the agenda is said to be endogenous. The difference between the existing work on issue-by-issue negotiation and ours is existing models study the strategic behaviour of agents by defining the agenda for all the issues either exogenously or endogenously. In contrast, we study the strategic behaviour of agents by defining an agenda that is partly exogenous and partly endogenous, and compare the outcome with the one that is generated if the agenda for all the issues is defined endogenously.

Our present work builds on the framework of [5, 6] which uses an endogenous agenda for all the issues. The players in this model use negotiation decision functions to generate offers and counter-offers [4]. The strategic behaviour of agents, which have time constraints and incomplete information, is studied for N negotiation issues.

Here, we study the effect of decomposing the set of N issues into k (for $1 \leq k \leq N$) subgroups and negotiating them sequentially, in stages, on the time of equilibrium agreement. The issues to be settled at each stage are determined exogenously. The negotiation protocol and equilibrium strategies of [5, 6] are then applied to each individual stage, and each agent's utility is determined for each value of k. The value of k that gives an agent the maximum utility is the optimal number of stages for that agent. The key result of our study is as follows. In some negotiation scenarios, the k-stage negotiation improves the utility of both agents, relative to the utility they get using the single stage negotiation protocol described in [5, 6]. However, identifying these scenarios requires each agent to know some private information about its opponent. Making an agent's private information known to the opponent can change the strategic behaviour of the opponent and result in decreased utility to the agent. It is therefore not feasible for agents to exchange their private information. We therefore extend the single stage negotiation protocol of [5, 6] by introducing a *mediator*. The mediator, that is assumed to be a neutral agent, obtains the required information from the two participants, and allows the players to improve their utilities by identifying such scenarios. We show that the extended k-stage negotiation protocol always gives each agent a utility that is no worse than the corresponding utility from single stage negotiation.

This paper makes two key advances to the state of the art in multi-issue negotiation. Firstly, it analyzes the strategic behaviour of agents by using an agenda that is partly exogenous and partly endogenous. Contrast this with the existing work that uses either an exogenous agenda or an endogenous agenda for all the N issues. Secondly, it presents an extended k-stage negotiation protocol that allows the players to improve their utilities, compared to the utilities from the single stage negotiation protocol.

The remainder of the paper is organised as follows. Section 2 gives an overview of the issue-by-issue bargaining model of [5, 6]. In Section 3 we describe our k-stage negotiation model. Section 4 describes the extended k-stage negotiation protocol. Section 5 discusses related work. Finally, in Section 6 we give some conclusions and make suggestions for future work.

2. THE SINGLE STAGE NEGOTIATION MODEL

We give a brief description of the multi-issue bargaining model of [5, 6], for which we determine the optimal agenda. A buyer, b, and a seller, s, that have unequal deadlines, bargain over the price of two[1] distinct issues, X and Y. Let a denote one of the two agents, i.e., $a \in \{b, s\}$, and $\hat{a}$ its opponent. Let T^a denote agent a's deadline for reaching agreement on both the issues. Agents have unequal deadlines and negotiation on all the issues must end by the earlier of the two deadlines. Otherwise, the agent with the earlier deadline quits and negotiation ends in a conflict; a situation which both agents prefer to avoid. We first describe the agents' information states, and then give an overview of the negotiation protocol and state the equilibrium outcomes.

2.1 Agents' Information State

We first consider the agents' information for a single issue, and then generalise it for more than one issue. For each issue, each agent has a reservation limit, a deadline, and a utility function. Thus b and s each have three parameters denoted $\langle RP^b, T^b, U^b \rangle$ and $\langle RP^s, T^s, U^s \rangle$ respectively. 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 defined as follows.

$$I^b = \langle RP^b, T^b, U^b, L_p^s, L_t^s \rangle$$

$$I^s = \langle RP^s, T^s, U^s, L_p^b, L_t^b \rangle$$

RP^b, T^b, and U^b are the information about b's own parameters and L_p^s and L_t^s are its beliefs about s. Similarly, RP^s, T^s, and U^s are s's own parameters and L_p^b and L_t^b are its beliefs about b. L_p^s and L_t^s are two lotteries that denote b's beliefs about s's deadline and reservation price. Agent b believes that s's reservation price is RP_l^s with probability α^b and RP_h^s with probability $1-\alpha^b$. It also believes that s's deadline is T_l^s with probability γ^b and T_h^s with probability $1-\gamma^b$. Thus

$$\begin{aligned} L_p^s &= [\alpha^b, RP_l^s; 1-\alpha^b, RP_h^s] \\ L_t^s &= [\gamma^b, T_l^s; 1-\gamma^b, T_h^s] \end{aligned}$$

Similarly L_p^b and L_t^b are two lotteries that denote s's beliefs about b's deadline and reservation price, where

$$\begin{aligned} L_p^b &= [\alpha^s, RP^b; 1-\alpha^s, RP^b] \\ L_t^b &= [\gamma^s, T_l^b; 1-\gamma^s, T_h^b] \end{aligned}$$

Thus agents have uncertain information about each other's deadline and reservation price. Moreover, the agents do not know their opponent's utility function or strategy.

When there is more than one issue to be negotiated (i.e., $N > 1$), RP^a in I^a represents a set of reservation prices for the N issues. The elements L_p^a and L_t^a now represents a set of N lotteries instead of a single lottery. The buyer and seller utilities are defined with the following two functions, U^b and U^s, that incorporate the effect of time discounting. Here, p_i is the price for issue i, t_i is the time

[1]This is a general model for negotiating N issues but for the sake of simplifying the discussion we consider only two issues.

of agreement for issue i, and δ_i^a is agent a's discounting factor for issue i.

$$U^b(p_1, \cdots, p_N, t_1, \cdots, t_N) = \sum_{i=1}^{N} (RP_i^b - p_i)(\delta_i^b)^{t_i}$$

$$U^s(p_1, \cdots, p_N, t_1, \cdots, t_N) = \sum_{i=1}^{N} (p_i - RP_i^s)(\delta_i^s)^{t_i}$$

When ($\delta_i^a > 1$) the agent gains utility with time and prefers a late agreement on issue i. When ($\delta_i^a < 1$) the agent loses utility with time and prefers an early agreement. Agents are said to have similar time preferences if both gain on time or both lose on time. Otherwise they have conflicting time preferences. Note that the discounting factors are different for different issues. This allows an agent to take a different attitude towards time for the different issues.

The utility functions can also be written in the following form.

$$U^b(p_1, \cdots, p_N, t_1, \cdots, t_N) = \sum_{i=1}^{N} (RP_i^b - p_i + c_i^b t_i)$$

$$U^s(p_1, \cdots, p_N, t_1, \cdots, t_N) = \sum_{i=1}^{N} (p_i - RP_i^s + c_i^s t_i)$$

Note that $\delta_i^a > 1$ in the first form corresponds to $c_i^a > 0$ in the second, and $\delta_i^a < 1$ corresponds to $c_i^a < 0$. Each agent's information state is its *private* knowledge that is not known to the opponent.

2.2 Negotiation Protocol

This is basically an alternating offers protocol. There are two types of offers. An offer on just one issue is referred to as a *single offer* and an offer on two issues 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 until agreement is reached on one of the issues. Thereafter, agents only make offers on the remaining issue. Negotiation ends when agreement is reached on both the issues or a deadline is reached. Let $X_{b\to s}^t$ denote the offer made on issue X by b to s at time t. The action, $A^s(t', X_{b\to s}^t, Y_{b\to s}^t)$, that agent s takes at time t' on a combined offer, 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.

The action A^b for b is defined analogously. A counter-offer for an issue is generated using the negotiation decision functions described in [13, 4]. 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, i.e., the interval $[RP^s, RP^b]$, 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 high 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 before negotiation begins. See [5, 6] for details on the equilibrium strategies of agents.

2.3 Equilibrium Outcomes

Here, T^b denotes b's deadline for reaching agreement on all the issues, and T^s denotes s's deadline for all the issues. For the above negotiation protocol, the strategic behaviour of each agent is studied on the basis of its own information state. In this model, the equilibrium outcome always results in an agreement on all the N issues. The equilibrium outcome on an issue depends on two factors: the relationship between the agent deadlines, and the agents' time preferences on that issue. Since the agents' utilities depend on the price, as well as time of agreement, the equilibrium outcome specifies both values. The equilibrium price for all the N issues depends on the relationship between agent deadlines, while the equilibrium time of agreement on an issue depends on the agents' time preferences for that particular issue and their deadlines. More specifically, in the equilibrium outcome, the price-surplus on all the issues goes to the agent with the longer deadline. The time of agreement on an issue depends on the agents' time preferences for that issue and the deadlines. When at least one agent prefers a late agreement on an issue i (i.e., $\delta_i^a > 1$ or $c_i^a > 0$ for at least one agent), then agreement on issue i takes place at the earlier deadline (denoted T). Otherwise the issue is agreed near the beginning of negotiation (denoted T'). Thus in some negotiation scenarios agreement on an issue occurs at the earlier deadline, while in others it occurs towards the beginning of negotiation.

3. THE K-STAGE NEGOTIATION MODEL

In the previous section, the order in which the N issues were bargained over, and agreements reached, is determined endogenously as part of the negotiation equilibrium. Also, the equilibrium time of agreement on each issue is either T (the earlier deadline) or T' (the beginning of negotiation), and depends on the agents' time preferences. More specifically, if at least one agent prefers a late agreement on an issue, then the issue is agreed at T, otherwise it is agreed at T'. Thus when agents have conflicting time preferences on an issue, the time of agreement for that issue is T. In other words, if ($\delta_i^a > 1$ and $\delta_i^{\hat{a}} < 1$), or ($c_i^a > 0$ and $c_i^{\hat{a}} < 0$) then agreement on issue i takes place at the latest possible time (T) and is therefore totally in favor of the patient agent. This holds true for all the N issues.

An agent can have different time preferences for different issues. Agent a gets maximum utility from time on the issues for which ($\delta_i^a > 1$) or ($c_i^a > 0$) and minimum utility from time on the issues for which ($\delta_i^a < 1$ and $\delta_i^{\hat{a}} > 1$) or ($c_i^a < 0$ and $c_i^{\hat{a}} > 0$). Our aim here is therefore to study the effect of decomposing the N issues into k stages and negotiating each stage sequentially (by applying the protocol and equilibrium strategies of Section 2 to each stage), on the time of equilibrium agreement, and how this in turn affects the agents' utilities. In other words, we determine if the cumulative utility from all the N issues can be improved for both the agents as a result of decomposition. We first give the formal definition of a k-stage decomposition and then determine the optimal agendas for both the agents.

3.1 Preliminary Definitions

We decompose the set of N issues into k equal subsets, and negotiate the subsets sequentially in k stages. The issues to be negotiated at each stage are defined exogenously (i.e., before negotiation begins). The issues that comprise each stage are negotiated by applying the protocol and equilibrium strategies of [5, 6]. In other words, the issues to be negotiated at each stage are determined exogenously, while the order in which the issues are settled at each stage is defined endogenously as part of the equilibrium outcomes of Section 2.3. For each stage, j ($1 \leq j \leq k$), we define a deadline

T_j^a for each agent. These deadlines are defined in such a way that the following two constraints are satisfied.

C1. For each agent, a, negotiation must end on all the stages by T^a. In other words, $T_k^a = T^a$ for both b and s, where T^a denotes agent a's deadline for reaching agreement on all the N issues, and $T_j^a < T^a$ for all values of j less than k.

C2. The relationship between agent deadlines for each of the k stages (i.e., the relationship between T_j^a and $T_j^{\hat{a}}$ for $1 \leq j \leq k$) has the same relationship as that between T^a and $T^{\hat{a}}$. In other words, the relation between agent deadlines is identical for all the stages and is independent of the value of k, i.e., decomposition preserves the relation between agent deadlines.

Note that in Section 2, $k = 1$. We begin by formally defining a *decomposition* of a set of negotiation issues and a *negotiation agenda*. We then determine the optimal agenda for each agent.

DEFINITION 1. *A k-stage decomposition, D_k, of a set S of N issues, is a sequence of k negotiation stages $D_k = A_1, A_2, \cdots, A_k$, where the stage A_i is defined as $A_i = \langle \frac{N}{k}, S_i, \frac{iT^b}{k} \frac{iT^s}{k} \rangle$. The number of issues in each stage is $\frac{N}{k}$, S_i is the set of $\frac{N}{k}$ issues selected randomly from S, $\frac{iT^a}{k}$ denotes agent a's deadline for the ith stage. S_i satisfies the following constraints.*

C3. $\cup_{i=1}^{k} S_i = S$

C4. $S_i \cap S_j = \emptyset$ *for all* i,j *such that* $i \neq j$

Constraint C3 ensures that each issue in the set S is settled at some stage, and constraint C4 ensures that no issue appears in more than one stage.

DEFINITION 2. *A k-stage agenda is a k-stage decomposition, together with a specification of the order in which the issues are settled at each stage. The order in which issues are settled at each stage is defined endogenously by the equilibrium strategies of [5, 6].*

DEFINITION 3. *The optimal agenda ($\mathbb{A}_o^a$) for agent a, for N negotiation issues, is the k_o^a-stage agenda which gives it maximum expected utility, i.e., $EU^a(D_o) > EU^a(D_x)$ for all $1 \leq x \leq N$ and $x \neq k_o^a$.*

We find the optimal agenda for each agent as follows. For each k-stage agenda, we find the expected utility (from all the N issues) to each agent. The value of k that gives an agent the maximum expected utility forms its optimal agenda.

Recall that in the previous section (for $k = 1$), the equilibrium price depends on the relationship between agent deadlines, while the equilibrium time of agreement depends on the agents' time preferences and the agent deadlines. Also, each agent has a single deadline, T^a, for completing negotiation on all the N issues. Since decomposition preserves the relation between agent deadlines (see constraint $C2$), and we use the endogenous agenda of [5, 6] for each stage, the equilibrium price for each of the N issues for a k-stage decomposition (where $k \neq 1$) remains identical to the equilibrium price given in Section 2 for $k = 1$. But since the N issues now have different deadlines (as opposed to Section 2 where each agent had a single deadline, T^a, for negotiating all the N issues), the equilibrium time of agreement for the N issues changes. In other words, the utility from price remains the same for both the agents, while their utility from time changes as a result of decomposition.

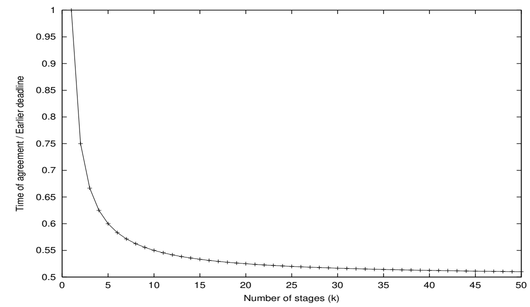

Figure 1: Effect of number of stages on the time of agreement when $T_1^i = T$

3.2 Optimal Agendas

The optimal agenda for both the agents is determined in the following steps.

1. For each value of k $(1 \leq k \leq N)$, obtain the k-stage decomposition.
2. For each value of k $(1 \leq k \leq N)$, determine the average equilibrium time of agreement on an issue.
3. Determine the optimal number of stages (k_o^a) for each agent.
4. Define all possible scenarios in which each agent can negotiate.
5. Determine the negotiation scenarios in which the optimal number of stages is identical for the two agents, i.e., $k_o^a = k_o^{\hat{a}}$.

If the optimal number of stages is identical for two the agents (and is not 1, as in the previous section), then it means that both the agents can benefit by negotiating the N issues in k sequential stages. Each of these steps is explained in detail below.

Step 1. For each value of k $(1 \leq k \leq N)$, we obtain the k-stage decomposition as per Definition 1. For example, if there are 10 issues to be negotiated between b and s, $T^b = 20$, and $T^s = 30$, then these 10 issues are split into, say 2 stages, each comprised of 5 issues. Agent b's deadline for the first stage is 10, and for the second stage it is 20. Agent s's deadline for the first stage is 15, and for the second stage is 30. The 2 stages are then negotiated sequentially using the protocol and equilibrium strategies of Section 2 for each individual stage.

Step 2. In order to determine the agent's utilities, we first determine the time of agreement for each issue. As stated earlier, the time of agreement on issue i depends on the agents' time preferences for that issue and the agents' deadlines. There are three possible scenarios in which negotiation can take place on an issue i.

S1. Both agents prefer a late agreement on issue i, i.e., ($\delta_i^a > 1$ and $\delta_i^{\hat{a}} > 1$) or ($c_i^a > 0$ and $c_i^{\hat{a}} > 0$).

S2. Both prefer an early agreement on issue i, i.e., ($\delta_i^a < 1$ and $\delta_i^{\hat{a}} < 1$) or ($c_i^a < 0$ and $c_i^{\hat{a}} < 0$).

S3. One of them prefers a late agreement while the other prefers an early agreement on issue i, i.e., ($\delta_i^a > 1$ and $\delta_i^{\hat{a}} < 1$) or ($c_i^a > 0$ and $c_i^{\hat{a}} < 0$).

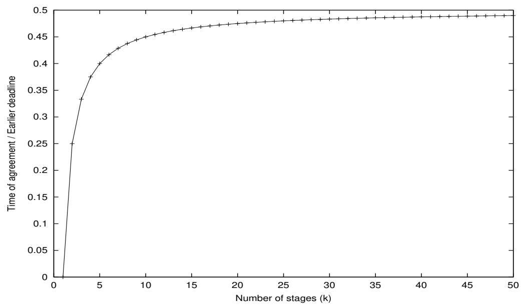

Figure 2: Effect of number of stages on the time of agreement when $T_1^i = T'$

On the basis of the equilibrium time of agreement on an issue for $k = 1$, we determine the equilibrium time of agreement on an issue for $k > 1$. This is done as follows.

Let T_k^i denote the equilibrium time of agreement on issue i when there are k stages. In Section 2.3 it was stated that the time of agreement T_1^i is T (the earlier deadline) in scenarios $S1$ and $S3$, and T' (the beginning of negotiation) in scenario $S2$. If the N issues are split into k stages, then issue i can become a part of any one of these k stages (since the issues for a particular stage are selected randomly from the set of N issues), say stage x. The new time of agreement for this issue depends on the time T_1^i at which it was agreed when $k = 1$. The two possible values for T_1^i are T and T'. Consider first the case where $T_1^i = T$. The new time of agreement for issue i, when $k > 1$, is the earlier deadline for stage x, i.e.,

$$T_k^i = \frac{xT}{k} \text{ where } 1 \leq x \leq k.$$

Since x can take any value between 1 and k, we compute T_k^i as the average equilibrium time of agreement over all the possible values of x. T_k^i therefore becomes

$$\begin{aligned} T_k^i &= \frac{1}{k}\sum_{x=1}^{k}\frac{xT}{k} \\ &= \frac{T(k+1)}{2k}. \end{aligned} \quad (1)$$

Figure 1 shows T_k^i for different values of k that lie in the range 1 and N, for $N = 50$. As seen in the figure, the average time of agreement on issue i is maximum at $k = 1$, and decreases as k increases. It reaches a minimum of $\frac{T}{2}$, when the number of stages is equal to the number of issues.

On the other hand, if $T_1^i = T'$ (i.e., the beginning of negotiation) then the new time of agreement for issue i, after decomposition, is the beginning of stage x, i.e.,

$$T_k^i = \frac{(x-1)T}{k} \qquad \text{where } 1 \leq x \leq k.$$

The average equilibrium time of agreement for issue i now becomes

$$\begin{aligned} T_k^i &= \frac{T}{k^2}\sum_{1}^{k-1} x \\ &= \frac{T(k-1)}{2k}. \end{aligned} \quad (2)$$

Figure 2 shows T_k^i for different values of k, that lie in the range 1 and N, for $N = 50$. As seen in the figure, the average equilibrium time of agreement, on issue i, is minimum at $k = 1$, and increases as k increases. It reaches a maximum of $\frac{T}{2}$, when the number of stages is equal to the number of issues.

Thus the average equilibrium time of agreement on an issue, for different values of k, can be obtained from Figure 1 (for scenarios $S1$ and $S3$) and Figure 2 (for scenario $S2$). For instance, the average equilibrium time of agreement on an issue, for $k = 5$, is $0.6T$ in scenarios $S1$ and $S3$, and $0.4T$ in scenario $S2$. For $k = N$, the average equilibrium time of agreement on an issue is $\frac{T}{2}$, in all three scenarios.

Step 3. Having obtained the average equilibrium time of agreement on an issue for all values of k between 1 and N, we determine the optimal number of stages for each agent. This is done as follows. For each value of k we compute the cumulative expected utility from all the N issues (EU_k^a), to each agent. The value of k for which EU_k^a is maximum is the optimal number of stages for agent a, and is denoted k_o^a. Note that although an agent's utility on an issue depends on both the price and time of agreement for that issue, we only need to determine the effect of decomposition on the utility from time. This is because (as stated earlier), the equilibrium price does not change as a result of decomposition; only the time of agreement does.

Let Y^a denote the number of issues on which agent a prefers a late agreement. This corresponds either to scenario S1 or S3. Let Z^a denote the number of issues on which agent a prefers an early agreement but the opponent prefers a late agreement, i.e., ($\delta^a < 1$ and $\delta^{\hat{a}} > 1$) or ($c^a < 0$ and $c^{\hat{a}} > 0$). This corresponds to scenario S3. The number of issues on which a and $\hat{a}$ prefer an early agreement, i.e., ($\delta^a < 1$ and $\delta^{\hat{a}} < 1$) or ($c^a < 0$ and $c^{\hat{a}} < 0$) is $N - Y^a - Z^a$. This corresponds to scenario S2. EU_k^a represents agent a's utility from time and is given by the following equation

$$EU_k^a = \sum_{i=1}^{Y^a}(c_i^a)T_1 + \sum_{j=1}^{Z^a}(c_j^a)T_1 + \sum_{m=1}^{N-Y^a-Z^a}(c_m^a)T_2. \quad (3)$$

In Equation 3, T_1 denotes the average equilibrium time of agreement on an issue in scenarios $S1$ or $S3$, and T_2 denotes the average equilibrium time of agreement in scenario $S2$. Note that for $k = N$, $T_1 = T_2 = \frac{T}{2}$, and for $k = 1$, $T_1 = T$ and $T_2 = T'$. Note also that $c_i^a > 0$, since the agent gains utility on time on Y^a issues, and c_j^a and c_m^a are both less than 0, since the agent loses utility on time on $N - Y^a$ issues. Using Equation 3, we compute the cumulative expected utility to each agent for all values of k between 1 and N.

EU_k^a depends on c_i^a, c_j^a, c_m^a, T_1, T_2, Y^a and Z^a. For each agent we fix the values of c_i^a, c_j^a and c_m^a to 1.2, -0.5 and -0.5 respectively (since $c_i^a > 0$ and $c_j^a < 0$ and $c_m^a < 0$). For a given value of k, we obtain T_1 and T_2 from Figure 1 and Figure 2 respectively. For each value of k, we then study the effect of Y^a and Z^a on EU_k^a. Consider first the effect of Y^a on the cumulative utility. When Y^a is large (i.e., the agent prefers a late agreement on most of the issues), the first term in Equation 3 contributes a larger share to EU_k^a than the remaining two terms. The first term has a maximum value at $k = 1$, since on each of the Y^a issues, this results in agreement at the latest possible time (i.e., T). For all values of Y^a greater than a particular value, say Y_c^a, EU_k^a was found to be maximum at $k = 1$. In our experiments Y_c^a was found to be 15. On the other hand, when $Y^a < Y_c^a$ (i.e., the agent prefers a late agreement on a small number of issues), the relationship between EU_1^a and EU_k^a (for $k > 1$) was found to depend on the value of Z^a. Consider the case where $Y^a < Y_c^a$. The relationship between EU_1^a and EU_k^a, for different values of k, is shown in Figure 3 for $Y^a = 1$. As seen in the figure, $EU_1^a = EU_N^a$ at $Z^a = 25$. Let Z_c^a denote this value

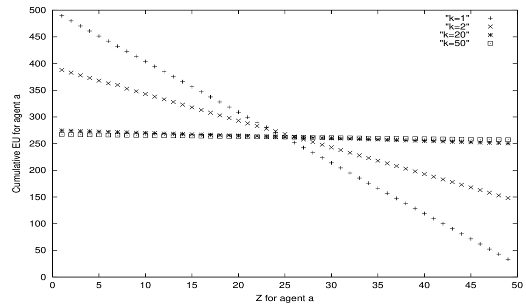

Figure 3: EU^a_k for different values of k when $Y^a < Y^a_c$

of Z^a at which $EU^a_1 = EU^a_N$.

For all values of Z^a such that $Z^a < Z^a_c$, $EU^a_1 > EU^a_k$, for all $k > 1$. On the other hand, for all values of Z^a such that $Z^a > Z^a_c$, $EU^a_N > EU^a_k$, for all values of $k < N$. Thus if $Z^a < Z^a_c$, then $k^a_o = 1$ and if $Z^a > Z^a_c$, $k^a_o = N$. In other words, when the number of issues on which ($c^a < 0$ and $c^{\hat{a}} > 0$) is low, then the single stage negotiation yeilds maxumum utility. On the other hand, when the number of issues on which ($c^a < 0$ and $c^{\hat{a}} > 0$) is high, then the decomposed N-stage negotiation yields maxumum utility. This result was found to hold true for all values of Y^a less than Y^a_c.

To summarize, the optimal mumber of stages (k^a_o) for agent a is 1 if $Y^a > Y^a_c$. On the other hand, if $Y^a < Y^a_c$, then k^a_o depends on the value of Z^a. For $Z^a < Z^a_c$, $k^a_o = 1$, and for $Z^a > Z^a_c$, $k^a_o = N$.

Step 4. On the basis of values for Y^a and Z^a, we define the following four negotiation scenarios.

S4. $Y^a > Y^a_c$

S5. $Y^a < Y^a_c$ and $Z^a < Z^a_c$

S6. $Y^a < Y^a_c$ and $Z^a > Z^a_c$

S7. $Y^a < Y^a_c$ and $Z^a = Z^a_c$

Each agent negotiates in one of the above four scenarios, which may be different for the two agents.

Step 5. On the basis of the value of k^a_o obtained in step 3, we determine the scenarios in which $k^a_o = k^{\hat{a}}_o$, i.e., the agents have identical values for the optimal number of stages. The different scenarios in which the two agents can interact, and the associated values of k^a_o are summarized in Table 1.

The following theorems give some results about the negotiation scenarios in which the agents have identical and conflicting values for the optimal number of stages.

THEOREM 1. *Agents a and $\hat{a}$ have conflicting values for the optimal number of stages in the following four scenarios.*

1. ($Y^a < Y^a_c$) *and* ($Z^a > Z^a_c$) *and* ($Y^{\hat{a}} > Y^{\hat{a}}_c$)
2. ($Y^a < Y^a_c$) *and* ($Z^a > Z^a_c$) *and* ($Y^{\hat{a}} < Y^{\hat{a}}_c$) *and* ($Z^{\hat{a}} < Z^{\hat{a}}_c$)
3. ($Y^a > Y^a_c$) *and* ($Y^{\hat{a}} < Y^{\hat{a}}_c$) *and* ($Z^{\hat{a}} > Z^{\hat{a}}_c$)
4. ($Y^a < Y^a_c$) *and* ($Z^a < Z^a_c$) *and* ($Y^{\hat{a}} < Y^{\hat{a}}_c$) *and* ($Z^{\hat{a}} > Z^{\hat{a}}_c$)

In all the remaining scenarios, the optimal number of stages is identical for the two agents, i.e., $k^a_o = k^{\hat{a}}_o$.

PROOF. *Consider scenario 1 first. This corresponds to scenario $S6$ for agent a and scenario $S4$ for agent $\hat{a}$ (i.e., entry in row 3, column 1 in Table 1). Here, $k^a_o = N$ and $k^{\hat{a}}_o = 1$, i.e., the agents have conflicting values for the optimal number of stages.*

Scenario 2 corresponds to scenario $S6$ for agent a and scenario $S5$ for agent $\hat{a}$ (i.e., entry in row 3, column 2 in Table 1). Here again $k^a_o = N$ and $k^{\hat{a}}_o = 1$, i.e., the agents have conflicting values for the optimal number of stages. Scenarios 3 and 4 correspond to (row 1, column 3) and (row 2, column 3) respectively, where $k^a_o = 1$ and $k^{\hat{a}}_o = N$. In all the remaining scenarios of Table 1, $k^a_o = k^{\hat{a}}_o$. □

THEOREM 2. *In the following negotiation scenarios, the N-stage negotiation gives maximum utility to both the agents.*

1. ($Y^a < Y^a_c$) *and* ($Z^a > Z^a_c$) *and* ($Y^{\hat{a}} < Y^{\hat{a}}_c$) *and* ($Z^{\hat{a}} > Z^{\hat{a}}_c$)
2. ($Y^a < Y^a_c$) *and* ($Z^a > Z^a_c$) *and* ($Y^{\hat{a}} < Y^{\hat{a}}_c$) *and* ($Z^{\hat{a}} = Z^{\hat{a}}_c$)
3. ($Y^a < Y^a_c$) *and* ($Z^a = Z^a_c$) *and* ($Y^{\hat{a}} < Y^{\hat{a}}_c$) *and* ($Z^{\hat{a}} > Z^{\hat{a}}_c$)
4. ($Y^a < Y^a_c$) *and* ($Z^a = Z^a_c$) *and* ($Y^{\hat{a}} < Y^{\hat{a}}_c$) *and* ($Z^{\hat{a}} = Z^{\hat{a}}_c$)

PROOF. *The first scenario corresponds to (row 3, column 3) of Table 1. In this scenario $k^a_o = k^{\hat{a}}_o = N$, i.e., the optimal number of stages is identical for the two agents. The remaining three scenarios correspond to (row 3, column 4), (row 4, column 3) and (row 4, column 4) of Table 1. In each of these 3 scenarios, the optimal number of stages for one of the agents is N while the other is indifferent between 1 and N. The N-stage negotiation thus gives a better utility to both the agents in all the above four scenarios.* □

To sum up, in some negotiation scenarios, both agents can improve their respective utilities by using the N-stage decomposition.

4. THE EXTENDED K-STAGE NEGOTIATION PROTOCOL

In the previous section, we showed that in some negotiation scenarios the N-stage negotiation yields maximum utility to both the agents. In order to identify such scenarios (listed in Theorem 2) the agents need to know the relationship between (Y^a and Y^a_c), ($Y^{\hat{a}}$ and $Y^{\hat{a}}_c$), (Z^a and Z^a_c), and ($Z^{\hat{a}}$ and $Z^{\hat{a}}_c$). However, recall that this information is not present in an agent's information state (see Section 2.1 for details on an agent's information state). Each agent plays the equilibrium strategy that is determined on the basis of its own information state. In order to allow the agents to get the maximum possible utility, we extend the negotiation protocol described in Section 2.2 by introducing a new agent called the *mediator*.

The role of the mediator is to facilitate the process of reaching an optimal agreement between the players by finding the optimal agenda. In order to determine the optimal agenda, the mediator first obtains each agents' time preference on each of the N issues. For each agent a, the mediator finds Z^a and sends it to agent a. Each agent then computes k^a_o and passes it on to the mediator. The mediator then determines the optimal agenda and sends it to both the agents. The agents then negotiate each stage as described in Section 2.

		$(Y^{\hat{a}} > Y_c^{\hat{a}})$ (1)	$(Y^{\hat{a}} < Y_c^{\hat{a}})$ and $(Z^{\hat{a}} < Z_c^{\hat{a}})$ (2)	$(Y^{\hat{a}} < Y_c^{\hat{a}})$ and $(Z^{\hat{a}} > Z_c^{\hat{a}})$ (3)	$(Y^{\hat{a}} < Y_c^{\hat{a}})$ and $(Z^{\hat{a}} = Z_c^{\hat{a}})$ (4)
1	$(Y^a > Y_c^a)$	$1/1$	$1/1$	$1/N$	$1/1, N$
2	$(Y^a < Y_c^a)$ and $(Z^a < Z_c^a)$	$1/1$	$1/1$	$1/N$	$1/1, N$
3	$(Y^a < Y_c^a)$ and $(Z^a > Z_c^a)$	$N/1$	$N/1$	N/N	$N/1, N$
4	$(Y^a < Y_c^a)$ and $(Z^a = Z_c^a)$	$1, N/1$	$1, N/1$	$1, N/N$	$1, N/1, N$

Table 1: Optimal number of stages in different negotiation scenarios. Rows represent the scenario for agent a and columns represent the scenario for agent $\hat{a}$. The entries indicate $k_o^a/k_o^{\hat{a}}$. A double entry $1, N$ indicates $EU_1^a = EU_N^a$

There are two reasons for having b and s go through the mediator instead of communicating directly. Firstly, if each agent lets the opponent know its time preference on each issue, then the information state of both the agents changes, i.e., it is no longer the information state described in Section 2.1 for which the equilibrium was computed. Each agent's strategic behaviour depends on its own information state and changing an agent's information state can therefore disturb the equilibrium of the game and, consequently, the outcomes listed in Section 2.3 may no longer be valid. The value of Z^a, that the mediator sends to the agents, only indicates the number of issues for which $c^a < 0$ and $c^{\hat{a}} > 0$ and not the actual issues for which $c^a < 0$ and $c^{\hat{a}} > 0$. So neither b nor s know the opponent's actual time preferences over the N issues, and the agents' strategic behaviour remains the same as in Section 2. Secondly, for the k-stage negotiation, the set of issues to be settled at each stage need to be selected randomly from the set of N issues. If either b or s is allowed to do this selection, it can select the issues for each stage on the basis of its time preferences (i.e., it could select the issues on which it prefers an early agreement for the earlier stages and the issues on which it prefers a late agreement for the later stages). We avoid this by having the mediator select the issues randomly for each stage.

The extended k-stage negotiation protocol is explained in more detail below.

1. Each agent sends its time preference on each issue to the mediator. Note that neither agent has the incentive to lie about this information. There are two possible values for an agent's time preference: it can either prefer a late agreement, or an early agreement on an issue. If it prefers a late agreement, then it can maximize the time of agreement only by sending its true time preference. Similarly, if it prefers a late agreement, then it can minimize the time of agreement only by sending its true time preference. Thus an agent is always better off by revealing its true time preference to the mediator.

2. The mediator computes Z^b and Z^s and sends these values to b and s respectively.

3. On the basis of Z^a, Y^a, Z_c^a and Y_c^a, each agent determines its optimal number of stages (k_o^a).

4. Each agent sends its value of k_o^a to the mediator.

5. If $k_o^a = k_o^{\hat{a}}$, then the mediator determines the optimal agenda. This is done as follows. There are two possible values for the optimal number of stages: 1 or N. If $k_o^a = 1$ for both the agents, then the N issues are negotiated in a single stage as described in Section 2. But if $k_o^a = N$ for both the agents, then the issues to be negotiated at each individual stage is determined randomly by the mediator. This forms the optimal agenda.

6. If $k_o^a = k_o^{\hat{a}}$, then the mediator sends the optimal agenda (obtained in the previous step) to both b and s. Otherwise it sends a conflict signal to both b and s.

7. If b and s receive an optimal agenda from the mediator, then they negotiate the stages sequentially, by applying the negotiation protocol and equilibrium strategies of Section 2 to each stage. On the other hand if they receive a conflict signal, they negotiate on all the N issues in a single stage as described Section 2.

THEOREM 3. *An agent's utility from the k-stage negotiation protocol is no worse than its utility from the single stage negotiation protocol.*

PROOF. *As shown in Table 1, there are only two possible values for the optimal number of stages: 1 or N. The agents have conflicting values in (row 1,column 3), (row 2,column 3), (row 3,column 1), and (row 3,column 2). In all the remaining scenarios they have identical values. Scenarios such as the one in (row 1,column 4) where agent $\hat{a}$ is indifferent between the two values are treated as identical since agent a's optimal number is one of these two values. Negotiation thus takes place in one of the following scenarios.*

1. *The optimal number of stages for the agents is identical and equal to 1.*

2. *The optimal number of stages for the agents is identical and equal to N.*

3. *The two agents have conflicting values for the optimal number of stages.*

In the first of the above three scenarios, each agent gets a utility that is equal to its utility from the single stage protocol. In the second scenario, the agents negotiate in N stages. The resulting utility to each agent is higher than its utility from the single stage protocol (see proof of Theorem 2). Finally, when the agents have conflicting values, they negotiate using the single stage protocol. Thus in scenarios 1 and 3, the utility to each agent from the extended k-stage protocol is equal to its utility from the single stage protocol. In scenario 2, each agent's utility from the extended protocol is higher than its utility from the single stage protocol. An agent's utility from the extended negotiation protocol is therefore no worse than its utility from the single stage negotiation protocol. □

5. RELATED WORK

Issue-by-issue negotiation has been explored in [8, 1, 9, 6, 5]. Fershtman [8] extends Rubinstein's complete information model [14] 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 [9] but it makes the agenda endogenous.

Bac and Raff [1] developed a model that has an endogenous agenda. They extended Rubinstein's model [15] 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 again *common knowledge*. Thus agents have asymmetric information about discounting factors.

While [8, 1, 2, 9] study the process of issue-by-issue negotiation in the absence of deadlines, the multi-issue model presented in [5, 6] analyses the effect of deadlines and discounting factors on issue-by-issue negotiation. However, the outcome generated by [5, 6], which uses an endogenous agenda for all the N issues, is not always Pareto optimal. We therefore determine if the utility to both agents can be improved by combining the exogenous agenda with the endogenous agenda. This is done by decomposing the N issues into k equal stages. The issues for each stage are determined exogenously, while the order in which issues are settled at each stage is determined endogenously. We then find the optimal agenda for each agent. Thus while existing work defines the agenda either exogenously for all the N issues, or endogenously for all the N issues, in our work the definition of the agenda is partly exogenous and partly endogenous.

Our study shows that in some scenarios, the utility to both agents can be improved by using the k-stage negotiation compared to the single stage negotiation. Moreover, our study also shows that the players do not have the ability to identify such scenarios, since they do not have complete information about all the negotiation parameters. We therefore presented a k-stage negotiation protocol, that allows agents to identify such scenarios, and thereby result in improved utility to both the agents.

6. CONCLUSIONS

This paper has studied the effect on the negotiation outcome of decomposing a set of N negotiation issues into k stages, and negotiating each stage sequentially. The set of issues to be negotiated at each stage is determined exogenously, while the order in which the issues are negotiated at each stage is determined endogenously as part of the bargaining equilibrium. Our study shows that the negotiation outcome changes with the value of k. Our study also shows that the optimal number of decompositions for an agent depends on the negotiation parameters. In some negotiation scenarios the optimal number of decompositions for the two agents differs, while in others it is identical. In other words, there exist negotiation scenarios in which the utility to both the agents can be improved by negotiating in stages compared to the utilities they get from single stage negotiation. However, since agents have incomplete knowledge of the negotiation parameters, they do not have the ability to identify these scenarios. We therefore presented an extended k-stage negotiation protocol, that allows agents to identify these scenarios through a mediator.

The proposed k-stage negotiation model thus offers two key advantages for agent mediated negotiation. Firstly, it is computationally simpler since it is basicaly an issue-by-issue model, and as such does not involve making tradeoffs between issues. Secondly, it results in improved utility to both the agents compared to the single stage model.

This paper studied the strategic behavior of agents by defining an agenda that was partly exognous and partly endogenous. In future, it would be interesting to explore the agents' strategic behaviour by separating the negotiation over the agenda from the negotiation over the issues.

Acknowledgements

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

7. 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] L. A. Busch and I. J. Horstman, Bargaining frictions, bargaining procedures and implied costs in multiple-issue bargaining, *Economica*, 64:669–680, 1997.

[3] L. A. Busch and I. J. Horstman, A Comment on issue-by-issue negotiations, *Games and Economic Behavior*, 19:144–148, 1997.

[4] P. Faratin and 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.

[5] S. S. Fatima and M. Wooldridge and N. R. Jennings, An agenda based framework for multi-issue negotiation, *Artificial Intelligence Journal*, (to appear).

[6] S. S. Fatima and M. Wooldridge and N. R. Jennings, Multi-issue negotiation under time constraints, In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'02)*, pages 143–150, Bologna, Italy, 2002.

[7] S. S. Fatima and M. Wooldridge and N. R. Jennings, Optimal negotiation strategies for agents with incomplete information, In J. J. Meyer and M. Tambe, editors, *Intelligent Agents VIII. Agent Theories, Architectures and Languages (LNAI Volume 2333)*, pages 377–392, Springer-Verlag: Berlin Germany, 2002.

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

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

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

[11] S. Kraus, *Strategic negotiation in multi-agent environments*, The MIT Press, Cambridge, Massachusetts, 2001.

[12] S. Kraus and J. Wilkenfeld and G. Zlotkin, Negotiation under time constraints, *Artificial Intelligence Journal*, 75(2):297–345, 1995.

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

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

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

[16] T. Sandholm and N. Vulkan, Bargaining with deadlines, In *Proceedings of the National Conference on Artificial Intelligence (AAAI'99)*, pages 44–51, Orlando, FL, 1999.

[17] T. Schelling, An essay on bargaining, *American Economic Review*, 46:281–306, 1956.