Optimal Negotiation of Multiple Issues in Incomplete Information Settings

Shaheen Fatima and Michael Wooldridge
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF, U.K.
{shaheen,mjw}@csc.liv.ac.uk

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

Abstract

This paper studies bilateral multi-issue negotiation between self-interested agents. The outcome of such encounters depends on two key factors: the ***agenda*** *(i.e., the set of issues under negotiation) and the negotiation* ***procedure*** *(i.e., whether the issues are discussed together or separately). Against this background, this paper analyses such negotiations by varying the agenda and negotiation procedure. This analysis is carried out in an incomplete information setting in which an agent knows its own negotiation parameters but has incomplete information about its opponent's parameters. We first determine the equilibrium strategies for two negotiation procedures: issue-by-issue and package deal. On the basis of these strategies we determine the negotiation outcome for all possible agenda–procedure combinations and the optimal agenda–procedure combination for each agent. We determine those conditions for which agents have identical preferences over the optimal agenda and procedure and those for which they do not, and for both conditions we show the optimal agenda and procedure.*

1. Introduction

Negotiation is a process by means of which agents communicate and compromise to reach mutually beneficial agreements [10, 8]. The simplest from of negotiation involves two agents and a single issue. The outcome of single issue negotiation depends on the agents' strategies, that, in turn, depend on the protocol and the information that agents have about the negotiation parameters (i.e., their reserve prices, their discount factors, their deadlines, and their utility functions) [9]. However, for multi-issue negotiation, the outcome also depends on two additional factors: the *agenda* and the *negotiation procedure* [14, 5, 16]. The agenda specifies what issues should be included in the negotiation[1], while the negotiation procedure specifies how the issues on the agenda will be settled. Generally speaking, there are two ways of negotiating multiple issues [6]. One approach is to discuss all the issues together as a *package deal*. The other approach is to settle each issue independently of all the other issues. This is called *issue-by-issue* negotiation. These two procedures, viz., issue-by-issue and package deal, yield different outcomes [6]. Consequently, the participants need to decide not only the agenda, but also the procedure they will use to settle the issues on the agenda.

In order to develop software agents that negotiate over multiple issues *optimally*, we need to identify general rules that can be used to determine the optimal agenda and procedure. Moreover, these rules need to be determined in the absence of complete information, since in most practical applications, agents do not have complete information about each other. To this end, this paper studies bilateral multi-issue negotiation between self interested agents in an incomplete information setting, and determines the optimal agenda and procedure for each agent in various scenarios.

Existing work on multi-issue negotiation has one of two key limitations. Either it is based on the complete information assumption [15], or else it focuses solely on the issue-by-issue procedure or solely on the agenda [5, 1, 7, 4]. This paper overcomes these limitations and thereby makes two important contributions to the state of the art in multi-issue negotiation. First, it treats both the agenda and the procedure as variable parameters and determines the optimal agenda and procedure for each agent in various scenarios. Second, it focuses on an incomplete information setting, in which neither agent has complete information about its opponent's preferences over possible outcomes.

The remainder of the paper is organised as follows. Section 2 describes the negotiation model. Section 3 obtains the optimal agenda and procedure for each agent. Section 4 discusses related work, and Section 5 concludes.

1 Note that the term agenda is generally used to mean the set of issues that will be negotiated and the order in which they will be negotiated. However, in this context, the term agenda refers only to the set of issues to be negotiated.

2. The negotiation model

Let a denote the buyer and $\bar{a}$ denote the seller (i.e., a's opponent). We use the term *information state* to refer to the information that an agent has about the negotiation parameters. Agent a's ($\bar{a}$'s) information state is denoted $\mathcal{I}^a$ ($\mathcal{I}^{\bar{a}}$). For our analysis, we consider the case where each agent has complete information about its own parameters but has incomplete information about its opponent's parameters. This is defined more precisely in the following subsections.

2.1. Single issue negotiation

Consider two agents negotiating over a single issue (A) that represents, say, the price of an object. Let T^a denote agent a's deadline by when it must have completed its negotiation. Also, let δ^a (that lies in the interval $(0, 1]$) denote the factor by which its utility gets discounted over time and R^a_A denote its reserve value. The deadline, discount factor, and reserve price for agent $\bar{a}$ are denoted $T^{\bar{a}}$, $\delta^{\bar{a}}$, and $R^{\bar{a}}$ respectively. Agents' utilities at price p and at time t, are defined as follows:

$$U^a_A(p,t) = \begin{cases} (R^a - p)(\delta^a)^t & \text{if } t \leq T^a \\ 0 & \text{if } t > T^a \end{cases}$$

$$U^{\bar{a}}_A(p,t) = \begin{cases} (p - R^{\bar{a}})(\delta^{\bar{a}})^t & \text{if } t \leq T^{\bar{a}} \\ 0 & \text{if } t > T^{\bar{a}} \end{cases}$$

An outcome is *individual rational* if it gives an agent a utility that is no less than its utility from the conflict outcome. Let U^a_{max} ($U^{\bar{a}}_{max}$) denote the maximum possible utility that agent a ($\bar{a}$) can get from an outcome that is individual rational to both agents. Agent a gets U^a_{max} for $p = RP^{\bar{a}}$ and $t = 0$. Analogously, $\bar{a}$ gets $U^{\bar{a}}_{max}$ for $p = RP^a$ and $t = 0$.

Assume that the agents use Rubinstein's alternating offers protocol [9]. At the beginning of negotiation, an agent makes an offer that gives it a high utility and then yields utility to its opponent by making offers that give it successively lower utility as negotiation progresses. This decrease is due to the competitive nature of the encounter and the fact that agents need to offer deals that are more likely to be accepted by their opponent if they are to come to an agreement. The generation of offers, i.e., how much an agent yields[2] depends on the negotiation parameters. Exactly which of the three actions (viz., accept, quit, or counter-offer) an agent takes and how it generates offers is defined by the agent's strategy.

An agent defines its strategy on the basis of the information it has about its opponent. For our analysis, we consider the case where each agent knows its own parameters, but has incomplete information about its opponent's parameters. The agents' information states are defined as follows:

$$\mathcal{I}^a = \{RP^a, T^a, \delta^a, U^a, RP^{\bar{a}}\}$$
$$\mathcal{I}^{\bar{a}} = \{RP^{\bar{a}}, T^{\bar{a}}, \delta^{\bar{a}}, U^{\bar{a}}, RP^a\}$$

Thus each agent knows its opponent's reserve price but does not know its opponent's deadline, discount factor, and utility function.

Agent a's strategy (denoted σ^a) is a mapping from the history of negotiation and agent a's information state to the action that it takes at time period t during negotiation. Thus action, $\mathtt{A}^a_t$, that agent a takes at time t is defined as:

$$\mathtt{A}^a_t = \sigma^a(h_t, \mathcal{I}^a, t) = \begin{cases} \text{Quit if } t > T^a \\ \text{Accept if } U^a_A(x^{\bar{a}}_t, t) \geq U^a_A(x^a_{t+1}, t+1) \\ \text{Offer } x^a_{t+1} \text{ at } t+1 \text{ otherwise} \end{cases}$$

where h_t is the history of negotiation and is a sequence of the form $x^a_1, x^{\bar{a}}_2, x^a_3, \ldots, x^{\bar{a}}_t$. x^a_t is the offer made by agent a at time t. Action $\mathtt{A}^{\bar{a}}_t$ is defined analogously.

Let $\sigma^a_e \times \sigma^{\bar{a}}_e$ denote the equilibrium[3] strategy profile. Let $\mathbf{x}^a_t$ denote the offer that agent a makes at time t in equilibrium (note that we use the upright boldface font for offers that are made at equilibrium). On the basis of σ^a_e we define agent a's *yield-factor*, y^a_t, at time t as:

$$y^a_t = \frac{U^a_{max} - U^a(\mathbf{x}^a_t, t)}{U^a_{max}} \quad (1)$$

where U^a_{max} is agent a's maximum possible utility. The yield factor for $\bar{a}$ is defined analogously.

2.2. Multi-issue negotiation

Consider the case where there are two issues (A and B). Let δ^a denote agent a's discount factor for both issues and T^a its deadline for completing negotiation on both issues. Similarly, agent $\bar{a}$'s deadline and discount factor are denoted as $T^{\bar{a}}$ and $\delta^{\bar{a}}$ respectively. The reserve prices for issue A and issue B are denoted as RP^a_A and RP^a_B for agent a, and $RP^{\bar{a}}_A$ and $RP^{\bar{a}}_B$ for agent $\bar{a}$. An agent's utility for agenda $\{A, B\}$ is the sum of its utilities from issue A and issue B. Agents' information states are defined as follows:

$$\mathcal{I}^a = \{RP^a_A, RP^a_B, T^a, \delta^a, U^a, RP^{\bar{a}}_A, RP^{\bar{a}}_B\}$$
$$\mathcal{I}^{\bar{a}} = \{RP^{\bar{a}}_A, RP^{\bar{a}}_B, T^{\bar{a}}, \delta^{\bar{a}}, U^{\bar{a}}, RP^a_A, RP^a_B\}$$

Definition 1 *An agenda, $\mathcal{A}$, is a set of issues that are included for negotiation. A negotiation procedure specifies how the set of issues on the agenda will be negotiated (i.e., if the issues will be negotiated one-by-one or whether all issues will be negotiated together as a package deal).*

2 See [2] for details on the different ways of yielding.

3 A number of game-theoretic models determine the equilibrium for the alternating offers protocol with different information states for incomplete information settings (see [12, 13] for details).

Multi-issue negotiation can be done using two procedures: *package deal* or *issue-by-issue* (as discussed in Section 1). The latter can in turn be done using two approaches. First, all issues can be discussed in parallel and independently of each other. We call this *parallel issue-by-issue* negotiation. The second approach is to negotiate the issues *sequentially*, one after another. For package deal, an offer includes a value for each issue under negotiation. Thus for two issues an offer is a *pair* of values, one for each issue. Agents are allowed to either accept a complete offer or reject a complete offer. This allows trade-offs to be made between issues. Furthermore, an agreement has to take place either on all the issues or none of them. On the other hand, for issue-by-issue negotiation each issue is dealt with separately and an agreement can take place either on a subset of issues or on all of them. For our present analysis we consider the parallel approach for issue-by-issue negotiation and compare it with the package deal[4].

Before obtaining the optimal agenda and procedure, we first describe the multi-issue negotiation protocol we use for our study and the agents' strategies. Consider the package deal procedure first. Assume that the agents use the same protocol as for single issue negotiation described in Section 2.1, but instead of making an offer on a single issue, an agent offers a set of pairs (a pair consists of an offer for issue A and issue B), all of which give it equal utility. This is because when there is more than one issue, an agent can make trade-offs across issues, resulting in a set of pairs, all of which give it equal utility. As an example, Figure 1(a) illustrates the utility frontiers for two issues: A and B. The segment AA' is the utility frontier for issue A and BB' that for issue B. The utility frontier for agenda $\{A, B\}$ is $A''B''C''D''$ (i.e., the sum of all possible utilities from issue A and issue B). For example, the points along JK (see Figure 1(a)) are pairs of values for issue A and issue B that give equal utility to agent $\bar{a}$ but different utilities to agent a. Out of all possible pairs along JK the only one that lies on the segment $A''B''C''$ is K, and this is Pareto-optimal. However, for our incomplete information setting, since an agent does not know its opponent's utility function, it does not know which of the possible pairs along JK is Pareto-optimal. Agent $\bar{a}$ therefore performs trade-offs[5] across A and B and offers a set of pairs that correspond to points along JK instead of a single pair.

Theorem 1 *For the package deal procedure, agent a's equilibrium strategy is as follows:*

$$\mathtt{A}^a_{t-1} = \begin{cases} \textit{Quit if } t-1 > T^a \\ \textit{Accept if offer received is better than counter offer} \\ \textit{Counter-offer } z^a_t \textit{ at } t \textit{ otherwise} \end{cases}$$

where z^a_t is a set of pairs of the form (x^a_{At}, x^a_{Bt}) that satisfy the following constraint, α:

$$U^a_{AB}(x^a_{At}, x^a_{At}, t) = U^a_{ABmax} - y^a_t \times U^a_{ABmax}$$

where $U^a_{ABmax} = U^a_{Amax} + U^a_{Bmax} = (RP^a_A + RP^a_B) - (RP^{\bar{a}}_A + RP^{\bar{a}}_B)$. U^a_{ABmax} is the maximum possible cumulative utility agent a can get from both issues. The equilibrium strategy for agent $\bar{a}$ is defined analogously.

Proof: Consider agent a. Since a has a single deadline for both issues, it quits if $t-1 > T^a$. The second action (i.e., accept) depends on the relative utilities of the offer agent a receives at time period $t-1$ to the counter offer it can send at time t. If the highest utility from the set of offers it receives is greater than the utility of its counter offer, it accepts. Otherwise it makes its counter offer. We show that, in equilibrium, agent a counter offers a set of pairs of the form (x^a_{At}, x^a_{Bt}) that satisfy constraint α, using proof by contradiction. Agent a's utility from such a pair is:

$$\begin{aligned} U^a_{AB}(x^a_{At}, x^a_{Bt}, t) &= (RP^a_A - x^a_{At} + RP^a_B - x^a_{Bt})(\delta)^t \\ &= (RP^a_{At} + RP^a_{Bt} - (x^a_{At} + x^a_{Bt}))(\delta)^t \end{aligned}$$

Assume that (x^a_{At}, x^a_{Bt}) is not an equilibrium offer. This implies that agent a can improve its utility (from the final outcome) by making a counteroffer at t that gives it a utility different from $U^a_{AB}(x^a_{At}, x^a_{Bt}, t)$. Let X^a_{At}, X^a_{Bt} denote such a pair. Assume that there is a single issue C with a reserve price of $RP^a_C = RP^a_A + RP^a_B$ to agent a and $RP^{\bar{a}}_C = RP^{\bar{a}}_A + RP^{\bar{a}}_B$ to $\bar{a}$. Also assume that the discount factor and deadline for issue C are δ^a and T^a respectively to agent a and $\delta^{\bar{a}}$ and $T^{\bar{a}}$ respectively to $\bar{a}$. In other words, at time t, the utility from offer x^a_{Ct} for the single issue C is analogous to the utility from offer X^a_{At}, X^a_{Bt} for issues A and B. Also, we know from Equation 1 that for single issue negotiation over C, agent a offers x^a_{Ct} at time t where x^a_{Ct} gives it a utility of: $U^a(x^a_{Ct}, t) = U^a_{Cmax} - y^a_t U^a_{Cmax}$ where U^a_{Cmax} is the maximum possible utility to a from issue C. We therefore have:

$$\begin{aligned} U^a_{Cmax} &= RP^a_C - RP^{\bar{a}}_C \\ &= (RP^a_A + RP^a_B) - (RP^{\bar{a}}_A + RP^{\bar{a}}_B) \quad (2) \end{aligned}$$

Since y^a_t is the yield factor in equilibrium for single issue negotiation, it is optimal for agent a to offer x^a_{Ct} at time t for the single issue C. This implies that the offer X^a_{At}, X^a_{Bt} (that corresponds to the offer x^a_{Ct}) is in equilibrium only if

4 Future work will deal with the sequential issue-by-issue approach.

5 Recall that although an agent does not know its opponent's utility, it knows its own utility function and can therefore perform trade-offs across issues.

$U^a(X^a_{At}, X^a_{Bt}, t) = U^a(x^a_{Ct}, t)$. We therefore have:

$$\begin{aligned} U^a(X^a_{At}, X^a_{Bt}, t) &= U^a(x^a_{Ct}, t) \\ &= U^a_{Cmax} - y^a_t U^a_{Cmax} \\ &= U^a_{ABmax} - y^a_t U^a_{ABmax} \quad (3) \\ &= U^a(x^a_{At}, x^a_{Bt}, t) \quad (4) \end{aligned}$$

This implies that a pair X^a_{At}, X^a_{Bt}, that results in agreement and gives agent a a utility greater than its utility from agreement that results from x^a_{At}, x^a_{Bt}, does not exist. All pairs of the form x^a_{At}, x^a_{Bt} that satisfy α therefore form an equilibrium for package deal. Agent $\bar{a}$'s equilibrium strategy can be obtained analogously. □

We now turn to the parallel issue-by-issue procedure. For this procedure, each issue is discussed in parallel using the single issue negotiation protocol described in Section 2.1.

Theorem 2 *For the parallel issue-by-issue procedure, agent a's equilibrium strategy is defined as follows:*

$$\mathtt{A}^a_{t-1} = \begin{cases} \textit{Quit if } t-1 > T^a \\ \textit{Accept issue } A \textit{ if } U^a(x^a_{A(t-1)}, t-1) \geq U^a(x^a_{At}, t) \\ \textit{Accept issue } B \textit{ if } U^a(x^a_{B(t-1)}, t-1) \geq U^a(x^a_{Bt}, t) \\ \textit{Counter-offer } x^a_{At} \textit{ if } A \textit{ is not agreed} \\ \textit{Counter-offer } x^a_{Bt} \textit{ if } B \textit{ is not agreed} \end{cases}$$

where x^a_{At} and x^a_{Bt} satisfy the following constraints: $x^a_{At} = U^a_{Amax} - y^a_t U^a_{Amax}$ and $x^a_{Bt} = U^a_{Bmax} - y^a_t U^a_{Bmax}$. Agent $\bar{a}$'s equilibrium strategy is defined analogously.

Proof: Unlike package deal, an agent here can accept a part of the offer it receives (i.e., only one of the two issues) or the complete offer (i.e., both issues). Assume that agent $\bar{a}$ plays its equilibrium strategy but a deviates and offers X^a_{At} instead of x^a_{At} for issue A. There are two possible relations between the corresponding utilities: $U^a(X^a_{At}, t) > U^a(x^a_{At}, t)$ or $U^a(X^a_{At}, t) < U^a(x^a_{At}, t)$. For competitive negotiation, if $U^a(X^a_{At}, t) > U^a(x^a_{At}, t)$ for agent a, then $U^{\bar{a}}(X^{\bar{a}}_{At}, t) < U^{\bar{a}}(x^{\bar{a}}_{At}, t)$ for $\bar{a}$. Moreover, since $\bar{a}$ plays its equilibrium strategy, it is possible that the offer $\bar{a}$ makes in the next time period gives it a utility higher than $U^{\bar{a}}(X^{\bar{a}}_{At}, t)$ but lower than $U^{\bar{a}}(x^{\bar{a}}_{At}, t)$. This implies that, at time t, $\bar{a}$ rejects the offer $X^{\bar{a}}_{At}$ but accepts $x^{\bar{a}}_{At}$. Furthermore, if t happens to be $\bar{a}$'s deadline (that a does not know) then $\bar{a}$ quits. Thus, if a offers $X^{\bar{a}}_{At}$, negotiation can end in a conflict but would otherwise have resulted in an agreement if a offered $x^{\bar{a}}_{At}$. This gives inferior utility to both agents. Thus it is not optimal for a to offer $X^{\bar{a}}_{At}$.

On the other hand, if $U^a(X^a_{At}, t) < U^a(x^a_{At}, t)$ for agent a we know that $U^{\bar{a}}(X^{\bar{a}}_{At}, t) > U^{\bar{a}}(x^{\bar{a}}_{At}, t)$ for agent $\bar{a}$. Since $\bar{a}$ follows its equilibrium strategy, it accepts $U^a(X^a_{At}, t)$, but might otherwise not do so if agent a offered x^a_{At}. This results in inferior utility to agent a relative to its equilibrium utility and a superior utility to agent $\bar{a}$. It is therefore not optimal for agent a to offer $U^a(X^a_{At}, t)$ in this case either. In other words, if agent $\bar{a}$ follows its equilibrium strategy, then agent a is better off by offering x^a_{At}. The same argument applies to the offer x^a_{Bt} for issue B. Agent a therefore offers x^a_{At} and x^a_{Bt} for issues A and B in equilibrium. Agent $\bar{a}$'s equilibrium strategy are obtained analogously. □

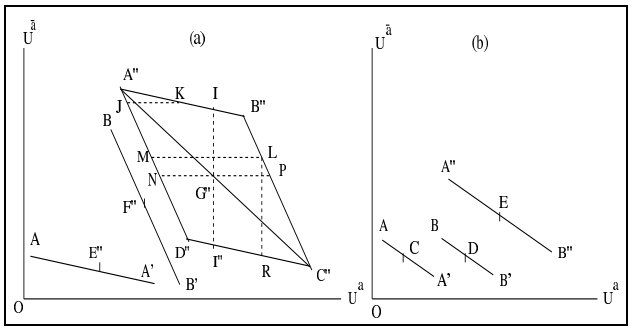

Figure 1. Agents' utilities in scenario S_1.

3. Optimal agenda and negotiation procedure

The negotiation outcome depends on both the agenda and the negotiation procedure. The optimal agenda and procedure are defined as follows.

Definition 2 *An agenda, $\mathcal{A}^a_o$ ($\mathcal{A}^{\bar{a}}_o$), is agent a's (agent $\bar{a}$'s)* optimal agenda *if, for either the issue-by-issue or the package deal procedure, it gives the agent the maximum utility between all possible agendas.*

An agent's utility from an agenda depends on the set of possible outcomes that both agents prefer over no deal (i.e., the zone of agreement for the agenda). On the basis of the zone of agreement for individual issues, we first list all possible scenarios in which negotiation can take place. We then find the optimal agenda, for each agent, for each possible scenario. The zone of agreement indicates outcomes that are *individual rational* to both agents. If agents' utilities are measured along the two axes, an issue has a zone of agreement if its utility frontier lies in quadrant Q_1. Each issue under negotiation may or may not have a zone of agreement. Thus for two issues we have four possible scenarios:

- S_1 Both issues have a zone of agreement.
- S_2 Only issue A has a zone of agreement.
- S_3 Only issue B has a zone of agreement.
- S_4 Neither issue A nor issue B has a zone of agreement.

3.1. Scenario S_1

In this scenario, the utility frontiers for both issues lie in quadrant Q_1. Figure 1(a) illustrates an example for this. Note that for the utility frontiers in this figure (and all the

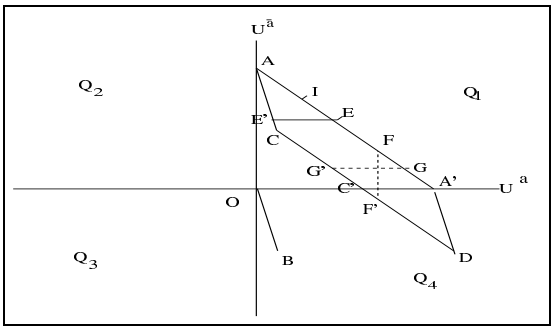

Figure 2. One sided gain to agent a in S_2.

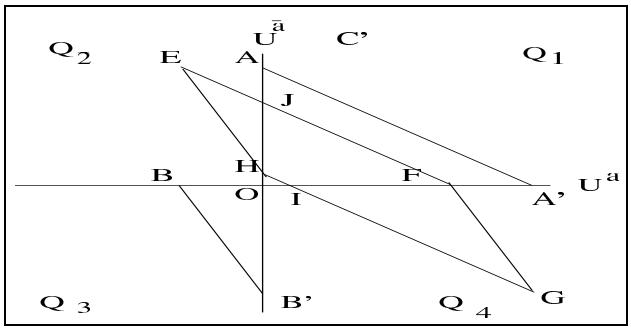

Figure 3. Agents' utilities in scenario S_2.

following ones) an increase in one agent's utility marks a decrease in the other agent's utility (i.e., negotiation is competitive). In all the figures, the origin represents the conflict outcome. Segment AA' represents the utility frontier for issue A and BB' that for issue B. When both issues are negotiated, the utility frontier is shown as $A''B''C''D''$ which represents the sum of utilities from issues A and B. Since the utility from agreement on each issue is higher than the conflict utility, each agent prefers to reach an agreement on both issues. In other words, when both issues have a zone of agreement, the optimal agenda includes both issues since each agent's cumulative utility from the two issues is greater than its utility from a single issue, irrespective of the negotiation procedure (i.e., $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \{A, B\}$).

When the optimal agenda includes more than one issue, these issues can be negotiated one-by-one or as a package deal (as per Section 2.2). These two procedures can generate different outcomes, and consequently give different utilities to the agents. The one that gives an agent a higher utility is its optimal procedure. For a given agenda, the outcomes generated by the two procedures depend on two factors: the corresponding equilibrium strategies and how the agents value the two issues (i.e., equally or differently). Agent a is said to value issue A more (less) than $\bar{a}$ if the increase in a's utility for a unit change in its allocation of the surplus for issue A is higher (lower) than the increase in $\bar{a}$'s utility for a unit change in $\bar{a}$'s allocation for the issue. Let ΔU_A^a denote the increase in agent a's utility due to a unit change in its allocation for issue A. Also, let $\mathcal{C}_A$ denote $\Delta U_A^a / \Delta U_A^{\bar{a}}$. We refer to $\mathcal{C}_A$ as the agents' *comparative interest* in issue A. Thus if $\Delta U_A^a / \Delta U_A^{\bar{a}} > 1$, agent a values issue A more than $\bar{a}$, and if it equals 1 then the agents value the issue equally. For example, in Figure 1(b), $(\mathcal{C}_A = \mathcal{C}_B)$.

Theorem 3 *If $\mathcal{C}_A \neq \mathcal{C}_B$, each agent's utility from the package deal is no worse than its utility from parallel issue-by-issue negotiation. If $\mathcal{C}_A = \mathcal{C}_B$, package deal and parallel issue by issue negotiation generate the same outcome.*

Proof: $\mathcal{C}_A \neq \mathcal{C}_B$: This situation is depicted in Figure 1(a). The two segments AA' and BB' denote agent a's and agent $\bar{a}$'s utilities for issues A and B respectively. The origin represents the conflict outcome. When the agenda includes both issues, the agents' combined utilities from the two issues lie in the region $A''B''C''D''$. Consider parallel issue-by-issue negotiation first. We know from Theorem 2 that in the equilibrium for this procedure, the offer agent a makes at time t is of the form $(\mathbf{x}_{At}^a, \mathbf{x}_{Bt}^a)$. For $\bar{a}$, it is $(\mathbf{x}_{At}^{\bar{a}}, \mathbf{x}_{Bt}^{\bar{a}})$. We also know that $\mathbf{x}_{At}^a(\mathbf{x}_{At}^{\bar{a}})$ is the offer that a ($\bar{a}$) would make if only the single issue A was being negotiated. The equilibrium time of agreement for single issue negotiation (denoted t_e) depends on the relation between the agents' deadlines and their discount factors (see [9] for details). Recall that each agent has a single deadline for both issues (T^a for agent a and $T^{\bar{a}}$ for $\bar{a}$). Furthermore, each agent has the same discount factor for both issues (i.e., δ^a and $\delta^{\bar{a}}$). Consequently, if the equilibrium time of agreement for single issue negotiation is t_e, the time of agreement for each issue, for parallel issue-by-issue negotiation is also t_e. Let E'' and F'' denote the equilibrium outcomes if each issue is negotiated using the single issue protocol of Section 2.1. Then the point G'' shows each agent's cumulative utility from both issues for parallel issue-by-issue negotiation.

Consider package deal, where an agent offers a set of pairs that give it equal utility. Recall that the offer that an agent, say a, makes at time t gives it a utility of $U_{ABmax}^a - y_t^a U_{ABmax}^a$. If at time t_e it is agent a's turn to make an offer that corresponds to G'', it offers a set of pairs that lie along II'', and if it is $\bar{a}$'s turn, it offers a set of pairs that lie along NP. At all time periods t' (where $t' < t_e$) each agent makes offers that give it a higher utility than G''. Let ML denote the offers made by $\bar{a}$ at time t' and LR denote those made by a. This results in an agreement at time t' at L. However, if the offer that $\bar{a}$ makes at time $t-1$ gives $\bar{a}$ a higher utility than I, and the offer that a makes at $t-1$ gives a a higher utility than P, agreement takes place at time t_e, either at I (if it is agent a's turn to offer at t_e) or at P (if it $\bar{a}$'s turn to offer at t_e). Thus, an agreement for package deal can occur anywhere along segment $IB''P$. Since all points on

$IB''P$ dominate G'' for both agents, the package deal gives each agent a utility that is no worse than its utility for the parallel issue-by-issue procedure. Thus if $\mathcal{C}_A \neq \mathcal{C}_B$, package deal gives each agent a utility that is no worse than its utility from issue-by-issue negotiation.

$\mathcal{C}_A = \mathcal{C}_B$: The two segments AA' and BB' denote a and $\bar{a}$'s utilities for issues A and B respectively (see Figure 1(b)). When the agenda includes both issues, the agents' combined utilities from the two issues lie along $A''B''$. Since no trade-offs are possible along $A''B''$, each agent always offers a single pair. An agreement for package deal therefore takes place at time t_e and gives each agent a utility that is the sum of its utilities from the equilibrium outcome for single issue negotiation on A and B. Thus if $\mathcal{C}_A = \mathcal{C}_B$, each agent gets equal utility from parallel issue-by-issue negotiation and package deal. □

3.2. Scenarios S_2 and S_3

In scenario S_2, issue A has a zone of agreement but issue B does not. Thus both agents prefer an agreement on issue A to no deal. For issue B, agents can have three possible preferences. First, a may prefer an agreement on issue B to no deal, while $\bar{a}$ prefers no deal to an agreement on issue B. Second, $\bar{a}$ may prefer an agreement on issue B to no deal while a prefers no deal to an agreement. Finally, neither agent may prefer an agreement on issue B to no deal. An agent's optimal agenda depends on these preferences.

Theorem 4 *In scenario S_2, there are two possible optimal agendas for the two agents:*

1. $\mathcal{A}_o^a = \{A, B\}$, $\mathcal{A}_o^{\bar{a}} = \{A\}$) *or* ($\mathcal{A}_o^a = \{A\}$, $\mathcal{A}_o^{\bar{a}} = \{A, B\}$)
2. $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \{A\}$

Proof: Since in scenario S_2 issue B does not have a zone of agreement, the only possible agendas are $\{A\}$ and $\{A, B\}$, one of which is optimal. The outcome for agenda $\{A, B\}$ can differ from the outcome for agenda $\{A\}$ in two ways. First, agenda $\{A, B\}$ can either increase the utility of one of the agents relative to agenda $\{A\}$ and reduce that of the other (i.e., result in one sided gain to one of them). Second, agenda $\{A, B\}$ can reduce the zone of agreement for issue A and result in decreased utility to both agents. We analyze each of these two cases.

1. Figure 2 illustrates the case where agenda $\{A, B\}$ improves the utility of a single agent. AA' represents the utility frontier for issue A and OB that for issue B. The origin is the conflict outcome. Here both agents prefer an agreement on issue A to no agreement, while only a prefers an agreement on issue B to no agreement. The utility frontier for agenda $\{A, B\}$ is $AA'DC$ and the corresponding zone of agreement is $AA'C'C$. Consider parallel issue-by-issue negotiation for agenda $\{A, B\}$. Let E denote the equilibrium outcome and t_e the time of agreement for single issue negotiation on A. Since issue B does not have a zone of agreement, separate negotiation on it results in conflict. E therefore represents the combined utility from both issues for the parallel issue-by-issue procedure.

   On the other hand, for package deal, the set of offers made by $\bar{a}$ at time t ($t < T^{\bar{a}}$) give it a utility of $U^{\bar{a}}_{ABmax} - U^{\bar{a}}_{ABmax} y^{\bar{a}}_t$. But $U^{\bar{a}}_{ABmax} = U^{\bar{a}}_{Amax} + U^{\bar{a}}_{Bmax}$ and $U^{\bar{a}}_{Bmax} = 0$. Thus $U^{\bar{a}}_{ABmax} = U^{\bar{a}}_{Amax}$. This implies that an offer that $\bar{a}$ makes at time t for single issue negotiation over A gives it the same utility as the offer it makes at time t for package deal. But for a, $U^a_{Amax} > 0$ and $U^a_{Bmax} > 0$. An offer that a would make at time t ($t < T^a$) for single issue negotiation over A would give it a utility of $U^a_{Amax} - U^a_{Amax} y^a_t$, while the offer it makes at time t for package deal gives it a utility of $U^a_{ABmax} - U^a_{ABmax} y^a_t$. Since $U^a_{ABmax} > U^a_{Amax}$, an offer that a makes at time t for package deal gives it a utility greater than its utility from its corresponding offer for the single issue A. Let I represent the offer made by a and E that made by $\bar{a}$, at equilibrium time t_e for the single issue A. From the above analysis it follows that the offers $\bar{a}$ makes at t_e for package deal lie along EE' and those that a makes lie to the right of I. Let FF' represent the offers made by a at t_e for the package deal. Since $E < F$ for agent a, an agreement does not take place at time t_e and negotiation proceeds to the next time period. Let $G'G$ denote the offers made by $\bar{a}$ at time t' (where $t' > t_e$). Since a's utility at G is greater than its utility at FF', a accepts the offer at G. Thus an agreement takes place at some point along FA'. Since FA' dominates E for a, and E dominates FA' for $\bar{a}$, the agents' optimal agendas are $\mathcal{A}_o^a = \{A, B\}$ and $\mathcal{A}_o^{\bar{a}} = \{A\}$. In general, if the utility frontier of issue B lies in quadrant Q_4, a gets one-sided gain from agenda $\{A, B\}$ relative to agenda $\{A\}$. In the same way it can be seen that if the utility frontier for issue B lies in quadrant Q_2, agenda $\{A, B\}$ improves $\bar{a}$'s utility and reduces that of a relative to agenda $\{A\}$. In this case, the optimal agendas are $\mathcal{A}_o^{\bar{a}} = \{A, B\}$ and $\mathcal{A}_o^a = \{A\}$.

2. Second, agenda $\{A, B\}$ can reduce the zone of agreement and decrease the utility of both agents relative to agenda $\{A\}$ (see Figure 3(a)). AA' and BB' are the utility frontiers for issues A and B. Both parties prefer an agreement on issue A to no agreement, while neither prefers an agreement on issue B to no agreement. The utility frontier for agenda $\{A, B\}$ is $EFGH$ and $JFIH$ is the corresponding zone of agreement. Since the points on AA' dominate those that lie in the region $JFIH$ for both agents, agenda $\{A, B\}$ decreases the

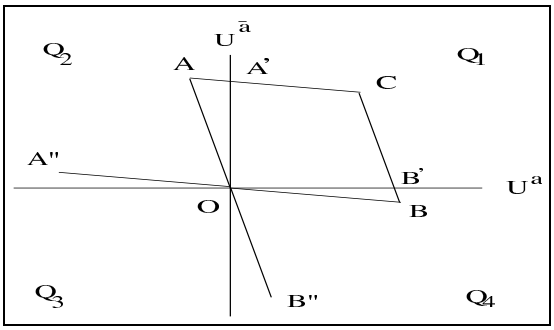

Figure 4. Agenda $\{A, B\}$ creates a zone of agreement in scenario S_4.

utility of both agents. Thus, if the utility frontier for issue B lies in quadrant Q_3, both agents prefer agenda $\{A\}$ to $\{A, B\}$. Thus $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \{A\}$. □

Thus in scenario S_2, agents may have similar as well as conflicting preferences over the optimal agenda. First, the optimal agenda for both agents is $\{A\}$ if agenda $\{A, B\}$ either destroys or reduces the zone of agreement giving inferior utility to both agents. Second, agenda $\{A, B\}$ can improve the utility of one agent and reduce that of the other, giving $\mathcal{A}_o^a = \{A, B\}$ and $\mathcal{A}_o^{\bar{a}} = \{A\}$.

As for scenario S_1, the optimal procedure for agenda $\{A, B\}$ is the package deal if $\mathcal{C}_A \neq \mathcal{C}_B$. If $\mathcal{C}_A = \mathcal{C}_B$, both procedures generate the same outcome. Scenario S_3 is analogous to scenario S_2.

3.3. Scenario S_4

In scenario S_4, neither issue has a zone of agreement and the utility frontier for neither issue lies in quadrant Q_1 (see Figure 4). OA represents the agents' utilities for issue A and OB those for issue B. The origin represents no deal. Agent a would not agree to any of the outcomes on OA, since none of them is better than no deal. Likewise, $\bar{a}$ would not agree to any of the outcomes on OB. Since settlement is preferable to no agreement only in quadrant Q_1, there are no settlement points for either issue. Each of the two issues, if discussed separately, would therefore result in conflict. However, negotiating the issues together is equivalent to taking the sum of utilities. Thus $A'CB'O$ is the zone of agreement for agenda $\{A, B\}$. Including both issues thus creates a zone of agreement and gives each agent a higher utility relative to no deal.

Theorem 5 *In scenario S_4, the optimal agenda is identical for the two agents and is either $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \{A, B\}$ or $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \emptyset$.*

Proof: Out of the four possible agendas, viz., $\emptyset$, $\{A\}$, $\{B\}$, or $\{A, B\}$, the two agendas $\{A\}$ and $\{B\}$ always result in conflict. However, if the sum of $\{A\}$ and $\{B\}$ creates a zone of agreement, both agents get a higher utility from agenda $\{A, B\}$ relative to the conflict outcome (see Figure 4). The optimal agenda is thus $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \{A, B\}$. But if the utility frontiers for issues A and B are $A''O$ and $B''O$ respectively, the sum of $\{A\}$ and $\{B\}$ does not create a zone of agreement. In this case, both agents prefer no deal to possible outcomes for agenda $\{A, B\}$ and $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \emptyset$. □

It is clear that if $\mathcal{A}_o^a = \mathcal{A}_o^{\bar{a}} = \{A, B\}$, then package deal is the optimal procedure for both agents.

These results are summarised in Tab. 1. As shown in the table, the optimal procedure depends on the agents' comparative interests, while the optimal agenda depends on the negotiation scenario. Furthermore, it is possible for agents to have identical preferences over the optimal agenda and procedure in each of the four scenarios.

	Condition	Optimal Agenda ($\mathcal{A}_o^a$)	Optimal procedure ($\mathcal{P}_o^a$)
S_1	×	$\mathcal{A}_o^a = \{A, B\}$ $\mathcal{A}_o^{\bar{a}} = \{A, B\}$	$\mathcal{P}_o^a = \mathcal{P}_o^{\bar{a}} = PD$ if $\mathcal{C}_A \neq \mathcal{C}_B$ $\mathcal{O}(IBI) = \mathcal{O}(PD)$ if $\mathcal{C}_A = \mathcal{C}_B$
S_2	a prefers agreement on issue B to no deal while $\bar{a}$ does not	$\mathcal{A}_o^{\bar{a}} = \{A\}$ $\mathcal{A}_o^a = \{A, B\}$	$\mathcal{P}_o^a = PD$ if $\mathcal{C}_A \neq \mathcal{C}_B$ $\mathcal{O}(IBI) = \mathcal{O}(PD)$ if $\mathcal{C}_A = \mathcal{C}_B$
	neither agent prefers agreement on issue B to no deal	$\mathcal{A}_o^a = \{A\}$ $\mathcal{A}_o^{\bar{a}} = \{A\}$	×
S_3	a prefers agreement on issue A to no deal while $\bar{a}$ does not	$\mathcal{A}_o^a = \{B\}$ $\mathcal{A}_o^{\bar{a}} = \{A, B\}$	$\mathcal{P}_o^{\bar{a}} = PD$ if $\mathcal{C}_A \neq \mathcal{C}_B$ $\mathcal{O}(IBI) = \mathcal{O}(PD)$ if $\mathcal{C}_A = \mathcal{C}_B$
	neither agent prefers agreement on issue A to no deal	$\mathcal{A}_o^{\bar{a}} = \{B\}$ $\mathcal{A}_o^a = \{B\}$	×
S_4	$\{A, B\}$ creates a zone of agreement	$\mathcal{A}_o^a = \{A, B\}$ $\mathcal{A}_o^{\bar{a}} = \{A, B\}$	$\mathcal{P}_o^a = \mathcal{P}_o^{\bar{a}} = PD$
	otherwise	$\mathcal{A}_o^a = \phi$ $\mathcal{A}_o^{\bar{a}} = \phi$	×

Table 1. Optimal agenda and negotiation procedure for the four scenarios. $\mathcal{O}(PD)$ denotes the outcome for the package deal and $\mathcal{O}(IBI)$ that for the parallel issue-by-issue procedure.

4. Related research

Existing game-theoretic models for multi-issue negotiation have two main limitations. First, they analyse the process of negotiation by assuming that agents have complete information. Second, they treat the agenda and procedure as

fixed parameters. For instance, Fershtman [5] extends Rubinstein's complete information model [11] for dividing a single pie to two pies. This model treats the agenda and the negotiation procedure as fixed parameters. It imposes an agenda exogenously, and studies the relation between the agenda and the outcome of the bargaining game for the sequential issue-by-issue procedure. It shows that the negotiation outcome, and, consequently, the agents' utilities, can be changed by changing the order in which issues are negotiated. Similar work in a complete information setting includes [7] but this makes the agenda endogenous. [15] treats the agenda as a variable parameter and shows how the negotiation outcome changes by changing the agenda. However, this work does not study how the negotiation procedure affects the outcome. Moreover, it is based on the complete information assumption.

Multi-issue negotiation models that deal with incomplete information include [1, 3, 4]. For instance, [1] develops a model that has an endogenous agenda. It extends Rubinstein's model for single pie bargaining with incomplete information [12] by adding a second pie. Similar work on endogenous agendas includes [3, 4]. However, while [1] considers uncertainty over deadlines, [3, 4] treats the opponent's deadline and reserve price as uncertain information. Furthermore, [3] treats each agent's information as its *private knowledge*, that is not known to its opponent. Although the above models deal with incomplete information, they treat both the agenda and the procedure as fixed parameters. In short, the work referenced above treats the issues to be negotiated as fixed and studies the relationship between the order (defined either exogenously or endogenously) in which the issues are discussed and the outcome of negotiation. In contrast, our work has the dual aim of determining the optimal agenda and procedure by treating both the agenda and procedure as variable parameters. Furthermore, we do not make the complete information assumption.

5. Conclusions and future work

This paper analysed the process of bilateral multi-issue negotiation by fixing the protocol and varying the agenda and negotiation procedure. We determined equilibrium strategies for two negotiation procedures: package deal and issue-by-issue. On the basis of these strategies, we determined the outcomes for all possible agenda–procedure combinations. Finally we showed the optimal agenda–procedure combination, for each agent, in various scenarios. Our study shows that although agents are self-interested, it is possible for them to have identical preferences over the optimal agenda and procedure in each possible scenario.

There are several interesting directions for future work. First, our analysis focussed on two issues. It is therefore important to extend it to more than two issues. Second, our study shows that although agents may have similar preferences over the agenda–procedure combination, they cannot recognise such scenarios due to lack of complete information. We therefore need to design a mechanism that allows them to do this and thereby get improved utilities.

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 J. J. Meyer and M. Tambe, editors, *Intelligent Agents VIII. Agent Theories, Architectures and Languages*, number 2333 in LNAI, pages 377–392. Springer Verlag, 2002.

[4] S. S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda based framework for multi-issue negotiation. *Artifi cial Intelligence Journal*, 152(1):1–45, 2004.

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

[6] C. Fershtman. A note on multi-issue two-sided bargaining: bilateral procedures. *Games and Economic Behavior*, 30:216–227, 2000.

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

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

[9] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

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

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

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

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

[14] J. C. Schelling. *The strategy of conflict*. Oxford University Press, 1960.

[15] J. K. Sebenius. Negotiation arithmetic: Adding and subtracting issues and parties. *International Organization*, 37(2):281–316, 1983.

[16] X. Zhang and V. Lesser. Multi-linked negotiation in multiagent systems. In *Proceedings of the First International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1207 – 1214, Bologna, Italy, 2002.