

On Efficient Procedures for Multi-Issue Negotiation

Shaheen S. Fatima¹ Michael Wooldridge¹ Nicholas R. Jennings²

¹ Department of Computer Science,
University of Liverpool, Liverpool L69 7ZF, U.K.
{S.S.Fatima, M.J.Wooldridge}@csc.liv.ac.uk

² 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 with deadlines. There are a number of procedures for negotiating the issues and each of these gives a different outcome. Thus, a key problem is to decide which one to use. Given this, we study the three main alternatives: the *package deal*, the *simultaneous procedure*, and the *sequential procedure*. First, we determine equilibria for the case where each agent is uncertain about its opponent's deadline. We then compare the outcomes for these procedures and determine the one that is optimal (in this case, the package deal is optimal for each party). We then compare the procedures in terms of their time complexity, the uniqueness and Pareto optimality of their solutions, and their time of agreement.

1 Introduction

Negotiation is a process that allows disputing agents to decide how to divide the gains from cooperation [5, 12]. Now, in practice, most negotiations involve multiple issues. However, for such encounters, the outcome depends on the *procedure* that is used [3]. Such procedures specify how the issues will be settled. Broadly speaking, there are three possibilities: (i) Discuss the issues together as a *package deal* (PD). This gives rise to the possibility of making tradeoffs across issues. (ii) Discuss the issues simultaneously, and independently of each other. This is called the *simultaneous procedure* (SIM). (iii) Discuss the issues one after another. This is called the *sequential procedure* (SEQ). Note that in the latter two cases, the issues are settled independently and so the agents cannot make tradeoffs.

As the three procedures yield different outcomes [10], a key problem for the agents is to decide what procedure they should use. Moreover, in many practical cases the agents have to decide this in the presence of time constraints and uncertain information. Given this, it is important to study the strategic behaviour of agents in such circumstances, and to determine what is the *optimal* procedure (i.e., the one that maximises expected utilities). To this end, this paper studies and compares the three main procedures for agents with deadlines and where each agent is uncertain about the other's deadline. We show that, for each agent, the PD is the best. We then compare the procedures in terms of four important attributes: their time complexity, whether their solutions are Pareto optimal, the uniqueness of their solutions, and their time of agreement.

Our analysis shows that, only the PD generates a Pareto optimal outcome, and that all three procedures have polynomial time complexity. In terms of the time of agreement, the PD and the SIM procedures are similar but the SEQ procedure is comparatively slower. Finally, we find the conditions for uniqueness of the solution.

There has been some formal comparison of different procedures to find the optimal one (see Section 5). However, all this work has two major limitations. First, it has focused on comparing procedures for negotiation without deadlines. But we believe deadlines are an important feature of most automated negotiations. Moreover, the strategic behaviour of agents with deadlines differs from that without. Second, it has only focused on finding the optimal procedure, but has not compared the solution properties of different procedures. Again, we believe this is a serious shortcoming that we rectify in this paper. Given this, our paper therefore makes a twofold contribution. First, we obtain the equilibrium for each procedure¹ when there are deadlines. Second, on the basis of this equilibrium, we provide the first comprehensive comparison of their solution properties (viz. time complexity, Pareto optimality, uniqueness, and time of agreement) and thereby allow agents to make a more informed choice about which procedure is most suitable in which circumstances.

The remainder of the paper is organised as follows. Section 2 introduces single-issue negotiation. Section 3 studies the three multi-issue procedures for the complete information scenario. Section 4 treats the agents' deadlines as uncertain. Section 5 discusses related work and Section 6 concludes.

2 Single-issue negotiation

We first give a reasonable standard model of single-issue negotiation and then move to the multi-issue case which is the main focus of this work. Two agents (*a* and *b*) negotiate over a single issue *i* using Rubinstein's alternating offers protocol [6]. Each agent has time constraints in the form of deadlines and discount factors. Since we focus on competitive scenarios with self-interested agents, we model negotiation with the 'split the pie game'. This complete information game is based on the split the pie game analysed in [2, 1]. The issue *i* is a 'pie' of size 1 and the agents want to find how to split it between themselves. The pie shrinks with time, and this shrinkage is represented by a discount factor denoted $0 < \delta_i \leq 1$ for both agents. At time $t = 1$, the size of the pie is 1, but at $t > 1$, the pie shrinks to δ_i^{t-1} . Let $n^a \in \mathbb{N}^+$ ($n^b \in \mathbb{N}^+$) denote agent *a*'s (*b*'s) deadline. If an agreement is not reached by an agent's deadline, then it quits and negotiation ends in a conflict. Both agents prefer an agreement to a conflict. Hence, negotiation must end by the earlier deadline $n = \min(n^a, n^b)$.

We denote the set of real numbers as \mathbb{R} and the set of real numbers in the interval $[0, 1]$ as \mathbb{R}_1 . Let $[x_i^t, y_i^t]$ denote the offer made at t where x_i^t and y_i^t denote *a*'s and *b*'s share respectively. Then, the set of possible offers is $\{[x_i^t, y_i^t] : x_i^t \geq 0, y_i^t \geq 0, \text{ and } x_i^t + y_i^t = \delta_i^{t-1}\}$ where $x_i^t \in \mathbb{R}_1$ and $y_i^t \in \mathbb{R}_1$. At time $t \leq n$, if *a* and *b* receive a share of x_i^t and y_i^t (where $x_i^t + y_i^t = \delta_i^{t-1}$), then their utilities are x_i^t and y_i^t

¹ Note that existing work has obtained equilibrium for negotiation with deadlines but only for the single issue case, and a special type of the SEQ procedure for multiple issues.

respectively. Agent a (b) gets zero utility if $t > n^a$ ($t > n^b$). Finally, the conflict utility is zero for both agents.

For this setting, the offers are determined as follows. Let a make an offer at $t = 1$. To begin, let the earlier deadline is $n = 1$. If b accepts at a 's offer at $t = 1$, the division occurs as agreed; if not, neither agent gets anything (since $n = 1$). Here, a is in a powerful position and is able to propose to keep 100 percent of the pie and give nothing to b ². Since $n = 1$, b accepts this offer and an agreement takes place at $t = 1$.

Now consider the case where the earlier deadline is $n = 2$. At $t = 1$, the size of the pie is 1 but it shrinks to δ_i at $t = 2$. In order to decide what to offer in the first round, a looks ahead to $t = 2$ and reasons backwards. It reasons that if negotiation proceeds to $t = 2$, b will take 100 percent of the shrunken pie by offering $[0, \delta_i]$ and leave nothing for a . Thus, at $t = 1$, if a offers b anything less than δ_i , b will reject the offer. Hence, at $t = 1$, a offers $[1 - \delta_i, \delta_i]$. Agent b accepts and an agreement occurs at $t = 1$.

In general, if the earlier deadline is n , a decides what to offer at $t = 1$ by looking ahead to $t = n$ and then reasoning backwards. This decision making leads a to offer $[\sum_{j=0}^{n-1}((-1)^j \delta_i^j), 1 - \sum_{j=0}^{n-1}((-1)^j \delta_i^j)]$ at $t = 1$. Agent b accepts and negotiation ends at $t = 1$. We now extend this single-issue model to the multi-issue case.

3 Multi-issue negotiation with complete information

As mentioned in Section 1, the existing literature does not analyse the multi-issue procedures for negotiation with deadlines³. Hence, we first analyse the complete information setting. Here, a and b negotiate over $m > 1$ issues. These issues are m distinct pies and the agents want to determine how to split each one. As before, each pie is of size 1. Let the discount factor for issue c where $1 \leq c \leq m$ be $0 < \delta_c \leq 1$. For each issue, let n^a (n^b) denote agent a 's (b 's) deadline. In the offer for time period t , a 's (b 's) share for each of the m issues is represented as an m element vector $x^t \in \mathbb{R}_1^m$ ($y^t \in \mathbb{R}_1^m$). Thus, if a 's share for issue c at time t is x_c^t , then b 's share is $y_c^t = (\delta_c^{t-1} - x_c^t)$. The shares for a and b are together represented as the package $[x^t, y^t]$.

An agent's *cumulative utility* from the package $[x^t, y^t]$ is the sum of its utilities for each of the m issues. Let $U^a : \mathbb{R}_1^m \times \mathbb{R}_1^m \times \mathbb{N}^+ \rightarrow \mathbb{R}$ and $U^b : \mathbb{R}_1^m \times \mathbb{R}_1^m \times \mathbb{N}^+ \rightarrow \mathbb{R}$ denote the cumulative utilities for a and b respectively at time $t \leq n$ where $U^a([x^t, y^t], t) = \sum_{c=1}^m k_c^a x_c^t$ and $U^b([x^t, y^t], t) = \sum_{c=1}^m k_c^b y_c^t$ where $k^a \in \mathbb{R}^m$ denotes an m element vector for a and $k^b \in \mathbb{R}^m$ that for b . These vectors indicate how the agents value different issues. For example, if $k_c^a > k_{c+1}^a$, then agent a values issue c more than issue $c+1$. Likewise for agent b . Each agent has complete information about all the negotiation parameters (i.e., n^a , n^b , m , k_c^a , k_c^b , and δ_c for $1 \leq c \leq m$). For this setting, we now obtain the equilibrium for the PD, the SIM, and the SEQ procedures.

The package deal procedure. For this procedure, the agents use the same protocol as

² It is possible that b may reject such a proposal. In practice, a will have to propose an offer that is just enough to induce b to accept. However, to keep the exposition simple, we assume that a can get the whole pie by making the 100 percent proposal.

³ The existing literature only analyses the case where each issue is discussed sequentially one after another (this is a special case of the procedures we study here). Section 5 gives details.

for the single-issue case (described in Section 2). However, an offer for the PD includes a proposal for each of the m issues. Agents are allowed to either accept a complete offer (i.e., all m issues) or reject a complete offer. An agreement can therefore take place either on all the m issues or none of them. As per single-issue negotiation, an agent decides what to offer by backward reasoning. However, since an offer for the PD includes a share for all the m issues, agents can now make tradeoffs across the issues in order to maximise their cumulative utilities. The function TRADEOFFA is agent a 's function for making tradeoffs, and is described in more detail in the proof of Theorem 1. The function TRADEOFFB for b can be defined analogously.

The equilibrium offer for issue c at time t is denoted as $[a_c^t, b_c^t]$, where a_c^t and b_c^t denote the shares for a and b . We denote the equilibrium package at time t as $[a^t, b^t]$ where $a^t \in \mathbb{R}_1^m$ ($b^t \in \mathbb{R}_1^m$) is an m element vector that denotes a 's (b 's) share for each of the m issues. Also, $\delta^{t-1} \in \mathbb{R}^m$ is an m element vector that represents the sizes of the m pies at time t . The symbol $\mathbf{0}$ denotes an m element vector of zeroes. For each pie, the sum of the agents' shares at time t is equal to the size of the pie at t (i.e., for $1 \leq t \leq n$, $a_c^t + b_c^t = \delta_c^{t-1}$). Finally, for time period $t \leq n$, we let $A(t)$ and $B(t)$ denote the equilibrium strategy for agents a and b respectively. Given this, Theorem 1 characterises the equilibrium for the PD.

Theorem 1. *The following strategies form a Nash equilibrium. For $t = n$ they are:*

$$A(n) = \begin{cases} \text{OFFER } [\mathbf{0}, \delta^{n-1}] & \text{if } a\text{'s turn} \\ \text{ACCEPT} & \text{if } b\text{'s turn} \end{cases} \quad (1)$$

$$B(n) = \begin{cases} \text{OFFER } [\mathbf{0}, \delta^{n-1}] & \text{if } b\text{'s turn} \\ \text{ACCEPT} & \text{if } a\text{'s turn} \end{cases} \quad (2)$$

For $t < n$, if $[x^t, y^t]$ denotes the offer made at time t , then the strategies are:

$$A(t) = \begin{cases} \text{OFFER } \text{TRADEOFFA}(\text{UB}(t)) & \text{if } a\text{'s turn} \\ \text{If } (U^a([x^t, y^t], t) \geq \text{UA}(t)) \text{ ACCEPT else REJECT} & \text{if } b\text{'s turn} \end{cases} \quad (3)$$

$$B(t) = \begin{cases} \text{OFFER } \text{TRADEOFFB}(\text{UA}(t)) & \text{if } b\text{'s turn} \\ \text{If } (U^b([x^t, y^t], t) \geq \text{UB}(t)) \text{ ACCEPT else REJECT} & \text{if } a\text{'s turn} \end{cases} \quad (4)$$

where $\text{UA}(t) = U^a([a^{t+1}, b^{t+1}], t+1)$ and $\text{UB}(t) = U^b([a^{t+1}, b^{t+1}], t+1)$. An agreement takes place at $t = 1$.

Proof. We look ahead to the last time period (i.e., $t = n$) and then reason backwards. If negotiation reaches the deadline (n), then the offering agent takes everything and its opponent gets nothing. Hence, we get Equations 1 and 2.

In all the preceding time periods ($t < n$), the offering agent proposes a package that gives its opponent a cumulative utility equal to what the opponent would get from its own equilibrium offer for the next time period. During time period t , either a or b could be the offering agent. Consider the case where a makes an offer at t . The package that a offers at t gives b a cumulative utility of $U^b([a^{t+1}, b^{t+1}], t+1)$. However, since there is more than one issue, there is more than one package that gives b this cumulative utility. Between these packages, a offers the one that maximises its own cumulative utility.

Thus, a 's tradeoff problem is to find a package $[a^t, b^t]$ that maximises $\sum_{c=1}^m k_c^a a_c^t$ such that $\sum_{c=1}^m (\delta_c^{t-1} - a_c^t) k_c^b = U^b([a^{t+1}, b^{t+1}], t+1)$ and $0 \leq a_c^t \leq 1$ for $1 \leq c \leq m$. This tradeoff problem is similar to the fractional knapsack problem [4, 14], the optimal solution for which can be found using the greedy approach (i.e., by filling the knapsack with items in their decreasing order of value per unit weight). The items in the knapsack problem are analogous to the issues in our case. The only difference is that the fractional knapsack problem starts with an empty knapsack and aims at filling it with items so as to maximise the cumulative value, while an agent's tradeoff problem can be viewed as starting with the agent having 100 per cent of all the issues and then aiming to give away portions of issues to the other agent so that the latter gets a given cumulative utility while the resulting loss in the former's utility is minimised. Hence, in order to perform tradeoffs, agent a considers k_c^a/k_c^b for $1 \leq c \leq m$ because k_c^a/k_c^b is the utility that a needs to give up in order to increase b 's utility by one. Since a wants to maximise its own utility and give b a utility of $U^b([a^{t+1}, b^{t+1}], t+1)$, it divides the m issues such that it gets the maximum possible share for those issues for which k_c^a/k_c^b is high and gives to b the maximum possible share for those issues for which k_c^a/k_c^b is low. Thus, a begins by giving b the maximum possible share for the issue with the lowest k_c^a/k_c^b . It then does the same for the issue with the next lowest k_c^a/k_c^b and repeats this process until b 's cumulative utility is $U^b([a^{t+1}, b^{t+1}], t+1)$. In this way, agent a performs tradeoffs with the TRADEOFFA($UB(t)$) function that uses the greedy approach described above. Thus we get Equation 3.

Analogously, if b offers at t , we get the equilibrium package of Equation 4. In this way, the first mover obtains the offer for $t = 1$ which its opponent accepts. \square

Theorem 2. For the PD, the time to find an equilibrium offer for $t = 1$ is $\mathcal{O}(mn)$.

Proof. The time to compute the equilibrium offer for $t = n$ is linear in the number of issues (see Equations 1 and 2). For $t < n$, the agents make tradeoffs. Recall from Theorem 1, that an agent's tradeoff problem is analogous to the fractional knapsack problem. Hence the time complexity TRADEOFFA (and TRADEOFFB) is $\mathcal{O}(m)$ (see [4, 14] for the complexity of the fractional knapsack problem). Tradeoffs are made in every time period from the $(n-1)$ th to the first. Hence the time complexity of finding an offer for $t = 1$ is $\mathcal{O}(mn)$. \square

Theorem 3. The PD has a unique equilibrium outcome if the following condition (C_1) is true:

$$C_1 : \quad \text{for all } i \text{ and } j, \text{ if } (i \neq j) \text{ then } (k_i^a/k_i^b \neq k_j^a/k_j^b)$$

Proof. Consider a time period $t < n$ and let a denote the offering agent. Recall from Theorem 1 that a splits the m issues in the increasing order of k_i^a/k_i^b . Thus, for a given i and j , if $k_i^a/k_i^b = k_j^a/k_j^b$, then agent a is indifferent between which of the two issues (i and j) it splits up first. For example, if $m = 2$, $n = 2$, $\delta = 0.5$, $k_1^a = 1$, $k_2^a = 2$, $k_1^b = 2$, and $k_2^b = 4$, then $k_1^a/k_1^b = k_2^a/k_2^b = 0.5$. If a is the offering agent at $t = 1$, it can offer $(1, 0)$ for issue 1 and $(1/4, 3/4)$ for issue 2. This gives a cumulative utility of 1.5 to a and 3 to b . Alternatively, a can offer $(0, 1)$ for issue 1 and $(3/4, 1/4)$ for issue 2 since this also results in the same cumulative utilities to a and b .

But if $k_i^a/k_i^b \neq k_j^a/k_j^b$, then a splits issue i first if $k_i^a/k_i^b < k_j^a/k_j^b$ and issue j first if $k_i^a/k_i^b > k_j^a/k_j^b$. Hence there is only one possible offer that a can make at any time

$t < n$. Likewise there is one possible offer that b can make at any time $t < n$. Since there is a unique offer for each time period, the equilibrium outcome is unique. \square

Theorem 4. *The PD generates a Pareto optimal outcome.*

Proof. As we consider competitive negotiations, for an individual issue c (where $1 \leq c \leq m$), an increase in one agent's utility results in a decrease in that of the other. However, for the PD procedure, an agent considers its cumulative utility from all m issues. Consequently, during the process of backward reasoning, at time $t < n$, the agent that makes tradeoffs maximises its own cumulative utility without lowering that of its opponent (with respect to what the opponent would offer in the next time period). Hence the equilibrium outcome for the PD is Pareto optimal. \square

The SIM procedure. Here the m issues are partitioned into $\mu > 1$ disjoint subsets. For $1 \leq c \leq \mu$, S_c denotes the c th partition, where $\cup_{c=1}^{\mu} S_c = \{1, \dots, m\}$. Negotiation for each partition starts at $t = 1$ and each partition is settled using the PD. Thus, for $\mu = m$, all m issues are settled simultaneously and independently of each other. At the other extreme, for $\mu = 1$, we have only one partition which is the PD procedure described earlier. Since the issues in each subset are settled using the PD, the equilibrium for each of these μ partitions is obtained from Theorem 1. Hence we get the following results.

First, an agreement for each issue occurs at $t = 1$. Since negotiation for each partition starts at $t = 1$ and an agreement for the PD occurs at $t = 1$ (see Theorem 1), an agreement for the SIM procedure (for each partition and hence each issue) occurs at $t = 1$. Second, if $|S_c|$ is the number of issues in S_c and n is the earlier deadline then the time to determine an equilibrium offer for $t = 1$ is $\sum_{c=1}^{\mu} \mathcal{O}(|S_c|n)$. Let M denote the size of the largest partition. Then, $\sum_{c=1}^{\mu} \mathcal{O}(|S_c|n) = \mathcal{O}(Mn)$. This is because the time to find the equilibrium offer for $t = 1$ for the PD (i.e., for $\mu = 1$) is $\mathcal{O}(mn)$ (see Theorem 2), so the time to compute equilibrium offer for $t = 1$ for the c th partition is $\mathcal{O}(|S_c|n)$. Hence, for all μ partitions, the time complexity is $\sum_{c=1}^{\mu} \mathcal{O}(|S_c|n)$. Third, it follows from Theorem 3 that the equilibrium outcome for the SIM procedure is unique if the condition C_1 is true for each of the μ partitions (irrespective of how the m issues are split into $\mu > 1$ partitions). Finally, as Theorem 5 shows, the SIM procedure does not always generate a Pareto optimal outcome.

Theorem 5. *The SIM procedure does not always generate a Pareto optimal outcome.*

Proof. We show this with a counter example. Let $n = 2$, $\delta = 0.5$, $m = 3$, $\mu = 2$, $S_1 = \{1, 2\}$, $S_2 = \{3\}$, $k_1^a = 1$, $k_2^a = 2$, $k_3^a = 3$, $k_1^b = 1$, $k_2^b = 0.5$, and $k_3^b = 0.25$. Let a denote the first mover. From Theorem 1, we know that in the equilibrium for partition S_1 , agent a gets a share of 0.25 for issue 1 and 1 for issue 2, and b gets a share of 0.75 for issue 1 and nothing for issue 2. For partition S_2 , each agent gets a share of 1/2. Thus, a 's cumulative utility from all the three issues is 3.75 and that of b is 0.875.

Now consider the case where all the three issues are discussed using the PD. Here, $\mu = 1$ and all other parameters remain the same. In the equilibrium outcome (i.e., the package $[(\frac{1}{8}, 1, 1), (\frac{7}{8}, 0, 0)]$), a gets a cumulative utility of 5.125 and b gets 0.875. This means that the procedure with $\mu = 2$ does not generate a Pareto optimal outcome. The reason for this is that the PD allows tradeoffs to be made across all the m issues while the simultaneous procedure only allows tradeoffs to be made across issues within each partition but not across partitions.

The SEQ procedure. The SEQ procedure differs from the SIM one in that the partitions are now negotiated sequentially, one after another. The issues within a subset are settled using the PD. Negotiation for the first partition starts at time $t = 1$. If negotiation for the c th (for $1 \leq c \leq \mu$) partition ends at t_c , then negotiation for the $(c + 1)$ th partition starts at time $t_c + 1$. Each agent gets its share for all the issues in a partition as soon as the partition is settled. Since the issues in each subset are settled using the PD, the equilibrium for each of these subsets is obtained from Theorem 1 by substituting the appropriate negotiation start times for each partition. Since negotiation for each partition ends in the same time period in which it starts, the time to settle all the m issues is μ . Note that the time complexity of the SEQ procedure is the same as the SIM one. Also, like the SIM procedure, the equilibrium for SEQ is not always Pareto optimal. Finally, the SEQ procedure has a unique outcome if the condition C_1 is true for all the partitions.

The optimal procedure. The procedure that gives a player the maximum utility is its optimal procedure. For the SEQ procedure the equilibrium outcome strongly depends on the negotiation *agenda* (i.e., the order in which the partitions are settled). There are two ways of defining the agenda [3]: *exogenously* (i.e., before the actual negotiation over the issues begins) or *endogenously* (the agents decide what issue they will settle next during the actual process of negotiation). The agenda that gives an agent the maximum utility is its optimal one [15]. Our objective here is not to determine the optimal agenda, but to consider a given agenda and compare the outcome for the SEQ procedure for the given agenda with the outcomes for the SIM and the PD procedures, in order to find the optimal one. The following theorem characterises this procedure.

Theorem 6. *Irrespective of how the m issues are split into $\mu > 1$ partitions, the PD is optimal for both parties.*

Proof. We first show that the PD is no worse than the SIM procedure. Consider the SIM procedure for $\mu > 1$. Since the difference between the procedure with $\mu = 1$ and that with $\mu > 1$ is that the former makes tradeoffs across all the m issues, while the latter does not, each agent's utility from the former is no worse than its utility from the latter.

We now show that for a given μ (where $\mu > 1$), for each agent, the outcome for the SIM procedure is better than that for the SEQ one (irrespective of the agenda for the SEQ procedure). We do this by considering each partition. Consider the partition $c = 1$. Since negotiation for the first partition starts at $t = 1$ for both SIM and SEQ procedures, the outcome for this partition is the same for $\mu = 1$ and $\mu > 1$. Hence, for the first partition, an agent gets equal utility from the two procedures. Now consider a partition $c > 1$. Let a denote the first mover for partition c (for $2 \leq c \leq \mu$) for both SIM and SEQ procedures. For the SIM procedure, negotiation for each partition starts at $t = 1$, and an agreement also occurs at $t = 1$. But, for the SEQ procedure, negotiation for the c th partition starts at $t = c$ and results in an agreement in the same time period. Since each pie shrinks with time, each agent's cumulative utility for the SIM procedure is greater than its cumulative utility for the SEQ one. Thus, for each agent, the PD is better than the SIM procedure, and the SIM procedure is better than the SEQ one.

We now extend the analysis to an incomplete information setting.

4 Multi-issue negotiation with uncertainty about deadlines

Here, there is uncertainty about the agents' deadlines. Both agents have a probability distribution over the possible values for n^a and n^b . Let $N \in \mathbb{N}^r$ denote a vector of r integers such that for $1 \leq i \leq r-1$, $N_i < N_{i+1}$. This vector represents the possible values for n^a and n^b (i.e., there are r types for a and r types for b). Let $P^a : \mathbb{N}^+ \rightarrow \mathbb{R}_1$ denote the discrete probability distribution function for n^a and $P^b : \mathbb{N}^+ \rightarrow \mathbb{R}_1$ that for n^b . The vector N and the functions P^a and P^b are common knowledge to the agents. Also, each agent knows its own type but not that of its opponent. In addition, each agent knows r , δ , k^a , k^b , and m . Since there are r possible types for each agent, we define r different cumulative utility functions for each of the two agents. If a is of type i (for $1 \leq i \leq r$) then its cumulative utility $U_i^a : \mathbb{R}_1^m \times \mathbb{R}_1^m \times \mathbb{N}^+ \rightarrow \mathbb{R}$ from the division specified by the package $[x^t, y^t]$ at time $t \leq N_i$ is $U_i^a([x^t, y^t], t) = \sum_{c=1}^m k_c^a x_c^t$ and zero if $t > N_i$. For b , $U_i^b = \sum_{c=1}^m k_c^b y_c^t$.

The PD procedure We know from Theorem 1, that the equilibrium outcome for the complete information setting depends on the earlier deadline n . In the present setting, since there is uncertainty about n , the equilibrium outcome now differs from that in Theorem 1. We first introduce some notation and then obtain the equilibrium.

Let $A(i, t)$ denote the equilibrium strategy for an agent a of type i at time t . Mutatis mutandis for $B(i, t)$. Let $[a^t, b^t]$ denote the package offered at t in equilibrium where $a^t + b^t = \delta^{t-1}$. Also, let $A(i, j, t)$ denote the equilibrium strategy for an agent a of type i for the time period t , assuming that b is of type j . Mutatis mutandis for $B(i, j, t)$.

Also, let $EUA(i, t)$ ($EUB(i, t)$) denote the cumulative utility that an agent a (b) of type i expects to get from b 's (a 's) equilibrium offer at time t . We let $EUA(i, j, t)$ denote agent a 's expected cumulative utility from its own equilibrium offer at time t if a is of type i , assuming that b is of type j ($EUB(i, j, t)$ is defined analogously). Note that the difference between $EUA(i, t)$ and $EUA(i, j, t)$ is that the former denotes a 's utility for the case where b is the offering agent at t , while the latter is a 's utility for the case where a is the offering agent at t . Likewise for $EUB(i, t)$ and $EUB(i, j, t)$.

Recall that in this setting, an agent only knows its own type but not that of its opponent. Since there are r possible types for each agent, there are r possible offers an agent can make at any time period (one offer corresponding to each possible type). Between these r offers, the one that gives an agent the maximum expected cumulative utility is its optimal offer. If the c th offer ($1 \leq c \leq r$) gives an agent the maximum expected cumulative utility, then we say that its *optimal choice* is c . For time period t , we let $OPTA(i, t)$ ($OPTB(i, t)$) denote the optimal choice for agent a (b) of type i .

Consider $t = N_r$. For this time period, for $1 \leq i \leq r$, we have the following (since N_r is the largest possible value for n):

$$EUA(i, N_r) = 0 \text{ and } EUB(i, N_r) = 0$$

$$EUA(i, j, N_r) = \begin{cases} 0 & \text{if } N_i < N_r \\ P^b(N_r) \times \left(\sum_{c=1}^m k_c^a \delta_c^{t-1} \right) & \text{if } N_i = N_r \end{cases}$$

$$\text{EUB}(i, j, N_r) = \begin{cases} 0 & \text{if } N_i < N_r \\ P^a(N_r) \times \left(\sum_{c=1}^m k_c^b \delta_c^{t-1} \right) & \text{if } N_i = N_r \end{cases}$$

Note that $\text{EUA}(i, j, N_r)$ and $\text{EUB}(i, j, N_r)$ do not depend on j because in the last time period, the offering agent gets 100 per cent of all the m pies. For $t < N_r$, we have:

$\text{EUA}(i, t) = \text{EUA}(i, \theta, t+1)$ and $\text{EUB}(i, t) = \text{EUB}(i, \lambda, t+1)$
where $\theta = \text{OPTA}(i, t+1)$ and $\lambda = \text{OPTB}(i, t+1)$.

$$\text{EUA}(i, j, t) = \begin{cases} 0 & \text{if } N_i < t \\ \sum_{e=1}^r \left(F^a(i, j, e, t) \times P^b(N_e) \right) & \text{if } N_i \geq t \end{cases}$$

$$\text{EUB}(i, j, t) = \begin{cases} 0 & \text{if } N_i < t \\ \sum_{e=1}^r \left(F^b(i, j, e, t) \times P^a(N_e) \right) & \text{if } N_i \geq t \end{cases}$$

The function F^a takes four parameters: i, j, e , and t , and returns the utility that an agent a of type i gets from offering the equilibrium package for time t , assuming that b is of type j but b is actually of type e . Obviously, b accepts a 's offer if $U_e^b(\text{A}(i, j, t), t) \geq \text{EUB}(e, \gamma, t+1)$ where $\gamma = \text{OPTB}(e, t+1)$. Hence, F^a is:

$$F^a(i, j, e, t) = \begin{cases} U_i^a(\text{A}(i, j, t)) & \text{if } U_e^b(\text{A}(i, j, t)) \geq \text{EUB}(e, \gamma, t+1) \\ \text{EUA}(i, t+1) & \text{otherwise} \end{cases}$$

where $\gamma = \text{OPTB}(e, t+1)$. The strategy $\text{A}(i, j, t)$ for $t = N_j$ is:

$$\text{A}(i, j, t) = \begin{cases} \text{OFFER } [\delta^{n-1}, \mathbf{0}] & \text{if } a \text{'}s \text{ turn} \\ \text{ACCEPT} & \text{otherwise} \end{cases}$$

and for all time periods $t < N_j$ it is:

$$\text{A}(i, j, t) = \begin{cases} \text{OFFER TRADEOFFA}(\text{EUB}(j, t)) & \text{if } a \text{'}s \text{ turn} \\ \text{if } U_i^a([x^t, y^t], t) \geq \text{EUA}(i, t) \text{ ACCEPT else REJECT} & \text{otherwise} \end{cases}$$

where $[x^t, y^t]$ is the package offered at t . Analogously, F^b is:

$$F^b(i, j, e, t) = \begin{cases} U_i^b(\text{B}(i, j, t)) & \text{if } U_e^a(\text{B}(i, j, t)) \geq \text{EUA}(e, \alpha, t+1) \\ \text{EUB}(i, t+1) & \text{otherwise} \end{cases}$$

where $\alpha = \text{OPTA}(e, t+1)$. The strategy $\text{B}(i, j, t)$ for $t = N_j$ is:

$$\text{B}(i, j, t) = \begin{cases} \text{OFFER } [\mathbf{0}, \delta^{n-1}] & \text{if } b \text{'}s \text{ turn} \\ \text{ACCEPT} & \text{otherwise} \end{cases}$$

and for all preceding time periods $t < N_j$ it is:

$$\text{B}(i, j, t) = \begin{cases} \text{OFFER TRADEOFFB}(\text{EUA}(j, t)) & \text{if } b \text{'}s \text{ turn} \\ \text{if } U_i^b([x^t, y^t], t) \geq \text{EUB}(i, t) \text{ ACCEPT else REJECT} & \text{otherwise} \end{cases}$$

Thus, the optimal choices for a and b are:

$$\text{OPTA}(i, t) = \arg \max_{j=1}^r \text{EUA}(i, j, t) \quad (5)$$

$$\text{OPTB}(i, t) = \arg \max_{j=1}^r \text{EUB}(i, j, t) \quad (6)$$

We compute the optimal choice for $t = 1$ by reasoning backwards from $t = N_r$. At $t = 1$, if an agent a of type i is the offering agent, then it offers the package that corresponds to b being of type $\text{OPTA}(i,1)$. Likewise, if an agent b of type i is the offering agent, then it offers the package that corresponds to a being of type $\text{OPTB}(i,1)$.

But since $\text{OPTA}(i,1)$ and $\text{OPTB}(i,1)$ are obtained under uncertainty, an agreement may or may not occur at $t = 1$. If it does not, then the agents update their beliefs as follows. Assume an agent a of type i makes an offer at $t = 1$. If this offer gets rejected, then it means that b is not of type $\text{OPTA}(i,1)$ and so a updates its beliefs about b using Bayes' rule (excluding passed deadlines and putting all the weight of the posterior distribution of a 's type over all N_i such that $i \neq \text{OPTA}(i,1)$). Now, on the basis of a 's offer at $t = 1$ (say $[a^1, b^1]$), b can infer the possible types for a . Thus, b also updates its beliefs using Bayes' rule (putting all the weight of the posterior distribution of a 's type over \mathcal{N} where $\mathcal{N} \subseteq N$ is the set of possible types for a that can offer $[a^1, b^1]$ in equilibrium). The belief update rules for the case where b offers at $t = 1$ are analogous to the case where a offers at $t = 1$. If the offer at $t = 1$ gets rejected, then negotiation goes to the next round. At $t = 2$, the offering agent (say an agent a of type i) finds $\text{OPTA}(i,2)$ with the updated beliefs. This process of updating beliefs and making offers continues until either an agreement is reached or one of the agents quits negotiation.

Theorem 7. *If $[x^t, y^t]$ denotes the offer made at time t , then for the PD procedure, for the time period $t \leq N_r$, the following strategies form a sequential equilibrium:*

$$A(i, t) = \begin{cases} QUIT & \text{if } t > N_i \\ OFFER \text{ TRADEOFFA}(\text{EUB}(\psi, t)) & \text{if } a \text{'s turn} \\ \text{If offer gets rejected UPDATE BELIEFS} & \\ RECEIVE OFFER and UPDATE BELIEFS & \text{if } b \text{'s turn} \\ \text{If } (U_i^a([x^t, y^t], t) \geq \text{EUA}(i, t)) \text{ ACCEPT else REJECT} & \end{cases} \quad (7)$$

$$B(i, t) = \begin{cases} QUIT & \text{if } t > N_i \\ OFFER \text{ TRADEOFFB}(\text{EUA}(\phi, t)) & \text{if } b \text{'s turn} \\ \text{If offer gets rejected UPDATE BELIEFS} & \\ RECEIVE OFFER and UPDATE BELIEFS & \text{if } a \text{'s turn} \\ \text{If } (U_i^b([x^t, y^t], t) \geq \text{EUB}(i, t)) \text{ ACCEPT else REJECT} & \end{cases} \quad (8)$$

for $1 \leq i \leq r$. Here, $\psi = \text{OPTA}(i, t)$ and $\phi = \text{OPTB}(i, t)$. Negotiation ends either in an agreement or a conflict. The earliest possible time of agreement is $t = 1$.

Proof. There are r possible values for the earlier deadline, and the vector N contains these possible values in ascending order. Hence, if $i < j$, then $\min(N_i, N_j)$ is N_i . To begin, consider the time period $t = 1$ and assume that an agent a of type i is the offering agent. There are r possible offers that a can make at t , one offer corresponding to each of the possible types for b (i.e., $A(i, j, 1)$ for $1 \leq j \leq r$). From these, a offers the one that gives it the maximum expected cumulative utility (i.e., the one with $j = \text{OPTA}(i, 1)$).

However, since $\text{OPTA}(i, 1)$ is computed under uncertainty (i.e., on the basis of expected utilities), an agreement may or may not take place at $t = 1$. If it does not, then negotiation proceeds as follows. Consider a time period t such that $1 \leq t < N_r$. Let

$[x^t, y^t]$ denote the offer made at time t . The agent that receives the offer (say a) updates its beliefs using Bayes' rule (excluding passed deadlines and putting all the weight of the posterior distribution of b 's type over \mathcal{N} where $\mathcal{N} \subseteq N$ is the set of possible types for b that can offer $[x^t, y^t]$ in equilibrium). If the proposed offer $([x^t, y^t])$ gets rejected, then the offering agent (say agent b of type i) updates its beliefs using Bayes' rule (putting all the weight of the posterior distribution of a 's type over all N_i such that $i \neq \text{OPTA}(i, 1)$). The belief update rules for the case where a offers at time t are analogous to the above rule. Hence we get Equations 7 and 8.

We now show that the beliefs specified above are consistent. During any time period $t < N_r$, suppose the strategy profile $(A(i, t), B(i, t))$ assigns probability $1 - \epsilon$ to the above specified posterior beliefs and probability ϵ to the rest of the support for the opponent's type. As $\epsilon \rightarrow 0$, the fully mixed strategy pair converges to (A, B) . Also, the beliefs generated by the fully mixed strategy pair converge to the beliefs described above. Given these beliefs, strategies A and B are sequentially rational.

We show the earliest possible time of agreement is $t = 1$ with an example: let $m = 2$, $\delta = 0.5$, $N_r = 2$, $r = 2$, $N = [1, 2]$, $k^a = [1, 2]$, $k^b = [2, 1]$, $P^a(1) = 0.1$, $P^a(2) = 0.9$, $P^b(1) = 0.9$, $P^b(2) = 0.1$. Let an agent a of type 1 (i.e., $n^a = 1$) be the offering agent at $t = 1$. Since $r = 2$, a can play two possible strategies at $t = 1$: one corresponding to the case where b is of type 1 and the other to the case where b is of type 2. For the former, a 's equilibrium offer at $t = 1$ is $[1, 0]$ for each issue. Hence $\text{EUA}(1, 1, 1) = 2.7$. For the latter case, a 's offer at $t = 1$ is $[0.325, 0.675]$ for the first issue and $[1, 0]$ for the second one. Hence $\text{EUA}(1, 2, 1) = 2.325$. Since $\text{EUA}(1, 1, 1) > \text{EUA}(1, 2, 1)$, $\text{OPTA}(1, 1) = 1$ and a plays the former strategy. Now if b is actually of type 1, then it accepts a 's offer. Thus, the earliest possible time of agreement is $t = 1$. But if b is of type 2, it rejects a 's offer since it can get a higher expected utility at $t = 2$. However, since a is of type 1, negotiation ends in a conflict.

If agent a 's offer at $t = 1$ gets rejected it knows that agent b is not of type $\text{OPTA}(i, 1)$. Thus the number of possible types for b is now reduced to $r - 1$. This happens every time a makes an offer that gets rejected. When negotiation reaches time period $t = 2r - 1$, there is only one possible type for b . An agreement therefore takes place at the latest by $t = 2r - 1$. However, if $n < 2r - 1$ then negotiation may end in a conflict.

Theorem 8. The time complexity of the PD procedure is $\mathcal{O}(mr^3T(N_r - \frac{T}{2}))$ where $T = \min(2r - 1, n)$.

Proof. Let a be the offering agent at $t = 1$ and let N_r be even (the proof for odd N_r is analogous). We begin with the last time period and then reason backwards. Since N_r is even and a starts at $t = 1$, it is b 's turn to offer in the last time period. For $t = N_r$, the time taken to find $\text{EUB}(i, j, t)$ (for a given i and j) is $\mathcal{O}(m)$ (see the definition of $\text{EUB}(i, j, t)$). Hence, the time taken to find $\text{EUB}(i, j, t)$ for all possible types of b (i.e., $1 \leq j \leq r$) is $\mathcal{O}(mr)$. Note that at this stage $\text{EUB}(j, t - 1)$ is known for $1 \leq j \leq r$.

Now consider $t = N_r - 1$. Since N_r is even, it is a 's turn to offer at $t = N_r - 1$. In order to find $A(i, t)$, we first need to find $\psi = \text{OPTA}(i, t)$. From the definition for $\text{OPTA}(i, t)$ we know that, for a given i , the time to find $\text{OPTA}(i, t)$ depends on the time to find $\text{EUA}(i, j, t)$ which in turn depends on the time to find $F^a(i, j, e, t)$. The time taken for $F^a(i, j, e, t)$ depends on the time taken for $A(i, j, t)$. For a given i and a given j , the time taken to find $A(i, j, t)$ is the time taken by the function TRADEOFFA. Since $\text{EUB}(j, t)$

is already known at time t , the time taken by TRADEOFFA is $\mathcal{O}(m)$ (see Theorem 2 for the complexity of TRADEOFFA). The time taken to find $F^a(i, j, e, t)$ is therefore $\mathcal{O}(m)$. Given this, the time to find $EUA(i, j, t)$ (for a given i and j) is $\mathcal{O}(mr)$. Hence, for a given i , the time to find $\psi = \text{OPTA}(i, t)$ is $\mathcal{O}(mr^2)$. Consequently, for a given i , the time to find $A(i, t)$ is $\mathcal{O}(mr^2)$. Recall that each agent knows only its own type and not that of its opponent. Hence we need to determine $A(i, t)$ for all possible types of a (i.e., for $1 \leq i \leq r$). This takes $\mathcal{O}(mr^3)$ time. Note that at this stage $EUA(i, j, t)$ is known for all possible values of i and j .

Now consider the time period $t = N_r - 2$ when it is b 's turn to offer. For $t = N_r - 2$ and a given i , the time to find $\text{OPTB}(i, t)$ is $\mathcal{O}(mr^2)$ and so the time to find $\text{OPTB}(i, t)$ for all possible types of b is $\mathcal{O}(mr^3)$. In the same way, the computation for each time period $t < N_r$ takes $\mathcal{O}(mr^3)$ time. Hence, the total time to find the equilibrium offer for $t = 1$ is $\mathcal{O}((N_r - 1)mr^3)$. However, as noted previously, an agreement may or may not occur at $t = 1$. If it does not, then the agents update their beliefs and find the equilibrium offer for $t = 2$. The time to compute the equilibrium offer for $t = 2$ is $\mathcal{O}((N_r - 2)mr^3)$. This process of updating beliefs and finding the equilibrium offer is repeated at most $T = \min(2r - 1, n)$ times (see the last paragraph of the proof for Theorem 7). Hence the time complexity of the PD is $\sum_{i=1}^T \mathcal{O}((N_r - i)mr^3) = \mathcal{O}(mr^3T(N_r - \frac{T}{2}))$.

Obviously, Theorems 3 and 4 extend to this scenario as well.

The SIM procedure. For the SIM procedure, the equilibrium for each partition is the same as that of Theorem 7. Consequently, the time complexity of computing an equilibrium offer is $\sum_{c=1}^{\mu} [\sum_{i=1}^T \mathcal{O}((N_r - i)|S_c|r^3)] = \mathcal{O}(Mr^3T(N_r - \frac{T}{2}))$. As before, M denotes the size of the largest partition. It is obvious that the condition for uniqueness is the same as that for the SIM procedure for the complete information case. Also, the outcome is not always Pareto optimal (see Theorem 5). Finally, for each partition, the earliest possible time of agreement is $t = 1$.

The SEQ procedure. For the SEQ procedure, the equilibrium outcome for the c th (for $1 \leq c \leq \mu$) partition is obtained from Theorem 7. The condition for uniqueness is the same as that for the SEQ procedure for the complete information case. The outcome is not always Pareto optimal (see Theorem 5). Also, the time complexity of the SEQ procedure is $\mathcal{O}(Mr^3T(N_r - \frac{T}{2}))$ (see Theorem 8). Finally, for the c th partition, the earliest possible time of agreement is $t_c = c$ (since the earliest possible time of agreement for the package deal is the first time period).

The optimal procedure. For each agent, the PD is optimal. The proof is analogous to the proof for Theorem 6 (except the fact that instead of actual utilities, we now use expected utilities).

5 Related work

A number of studies have analysed different procedures for multi-issue negotiation. For instance, Fershtman [3] extended Rubinstein's model [2] for splitting a single pie

	Package deal	Simultaneous	Sequential
Time of agreement (t_c)	Earliest possible time for the c th issue $t_c = 1$ for $1 \leq c \leq m$	Earliest possible time for the c th issue $t_c = 1$ for $1 \leq c \leq m$	Earliest possible time for the c th partition $t_c = c$ for $1 \leq c \leq \mu$
Time to compute equilibrium	$\mathcal{O}(mr^3T(N_r - \frac{T}{2}))$	$\mathcal{O}(Mr^3T(N_r - \frac{T}{2}))$	$\mathcal{O}(Mr^3T(N_r - \frac{T}{2}))$
Pareto optimal?	Yes	No	No
Conditions for uniqueness	if C_1 is true	if C_1 is true for every partition	if C_1 is true for every partition

Table 1. Outcomes for the incomplete information setting – m is the total number of issues and M is the number of issues in the largest partition (for a definition of C_1 see Theorem 3).

to SEQ negotiation for two pies. This model assumes complete information, imposes an agenda exogenously, and studies the relation between the agenda and the outcome of the SEQ bargaining game. On the other hand, [11, 13, 7] study negotiations with an endogenous agenda. For instance, [11] studies PD, SIM, and SEQ negotiation by assuming complete information. Furthermore, the agents are assumed to have discount factors but no deadlines. The main result of this work is that the PD is the optimal procedure and that for each procedure there exist multiple equilibria. [13] extends this work by finding conditions under which the equilibrium is unique. [7] developed an *asymmetric information* model for two issues and studied the PD and the SEQ procedure. A slightly different approach was taken in [8] by adding a preliminary period in which agents bargain over an agenda first and then settle the issues using this agenda. However, in [7] and [8] the players have discount factors but no deadlines. In summary, the above work differs from ours in that we consider both discount factors and deadlines, whereas previous work only considers discount factors and no deadlines⁴. Negotiation with deadlines was studied in [9] but only for a single issue. Also, the existing literature does not compare the different procedures in terms of a comprehensive list of their attributes (viz. time complexity, Pareto optimality, uniqueness, and time of agreement). Our comparative study of these attributes allows a more informed choice to be made about which procedure is most suitable in which circumstances.

6 Conclusions and future work

This paper analysed the three key procedures for bilateral multi-issue negotiation between self-interested agents: the PD, the SIM, and the SEQ procedures. Our results (see Table 1) show that the PD is better than the other two because it is the optimal procedure for both agents, it is the only one to generate a Pareto optimal outcome, and it achieves these with polynomial time complexity (as the other two procedures). With regard to the time of agreement, the PD and the SIM procedures are similar in that, for the complete information setting, both procedures result in an agreement at $t = 1$ for all the issues. Also, when there is uncertainty about deadlines, the earliest possible time of agreement is the same for both procedures. But the SEQ procedure is slower in terms

⁴ [15] only determines the optimal agenda for SEQ negotiation (with a single issue in each partition) with deadlines.

of the time of agreement. Finally, all the three procedures have a unique outcome under certain conditions.

In future, we will extend our symmetric information analysis by studying asymmetric information settings. Also, in this work, we modelled the players' time preferences in the form of discount factors. However, it has been shown that the outcome for negotiation with discount factors can differ from that for fixed time costs [8]. Therefore, it will be interesting to extend our analysis to negotiations with fixed time costs.

Acknowledgements

We are grateful to Sarit Kraus for her detailed comments on earlier versions of this paper.

References

1. I. Stahl, *Bargaining Theory*, Economics Research Institute, Stockholm School of Economics, Stockholm, 1972.
2. A. Rubinstein, Perfect equilibrium in a bargaining model, *Econometrica*, 50(1):97–109, January, 1982.
3. C. Fershtman, The importance of the agenda in bargaining, *Games and Economic Behavior*, 2:224–238, 1990.
4. S. Martello and P. Toth, *Knapsack problems: Algorithms and computer implementations*, (Chapter 2), John Wiley and Sons, 1990.
5. J. S. Rosenschein and G. Zlotkin, *Rules of Encounter*, The MIT Press, 1994.
6. M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, The MIT Press, 1994.
7. M. Bac and H. Raff, Issue-by-issue negotiations: the role of information and time preference, *Games and Economic Behavior*, 13:125–134, 1996.
8. L. A. Busch and I. J. Horstman, Bargaining frictions, bargaining procedures and implied costs in multiple-issue bargaining, *Economica*, 64:669–680, 1997.
9. 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.
10. C. Fershtman, A note on multi-issue two-sided bargaining: bilateral procedures, *Games and Economic Behavior*, 30, 216–227, 2000.
11. R. Inderst, Multi-issue bargaining with endogenous agenda, *Games and Economic Behavior*, 30: 64–82, 2000.
12. S. Kraus, *Strategic negotiation in multi-agent environments*, The MIT Press, Cambridge, Massachusetts, 2001.
13. Y. In and R. Serrano, Agenda restrictions in multi-issue bargaining (II): unrestricted agendas, *Economics Letters*, 79:325–331, 2003.
14. T. H. Cormen and C. E. Leiserson and R. L. Rivest and C. Stein, *An introduction to algorithms*, The MIT Press, Cambridge, Massachusetts, 2003.
15. S. S. Fatima and M. Wooldridge and N. R. Jennings, An agenda based framework for multi-issue negotiation, *Artificial Intelligence Journal*, 152(1):1–45, 2004.