

An Analysis of Feasible Solutions for Multi-Issue Negotiation Involving Nonlinear Utility Functions

Shaheen Fatima
Department of
Computer Science
Loughborough University
Loughborough LE11 3TU, UK.
s.s.fatima@lboro.ac.uk

Michael Wooldridge
Department of
Computer Science
University of Liverpool
Liverpool L69 3BX, UK.
mjw@csc.liv.ac.uk

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

ABSTRACT

This paper analyzes bilateral multi-issue negotiation between self-interested agents. Specifically, we consider the case where issues are *divisible*, there are time constraints in the form of *deadlines* and *discount factors*, and the agents have different preferences over the issues. Given these differing preferences, it is possible to reach Pareto-optimal agreements by negotiating all the issues together using a *package deal procedure* (PDP). However, finding equilibrium strategies for this procedure is not always computationally easy. In particular, if the agents' utility functions are *nonlinear*, then equilibrium strategies may be hard to compute. In order to overcome this complexity, we explore two different solutions. The first is to use the PDP for linear approximations of the given nonlinear utilities. The second solution is to use a *simultaneous procedure* (SP) where the issues are discussed in parallel but independently of each other. We then compare these two solutions both in terms of their computational properties (i.e., time complexity of computing an approximate equilibrium and the associated error of approximation) and their economic properties (i.e., the agents' utilities and social welfare of the resulting outcome). By doing so, we show that an approximate equilibrium for the PDP and the SP can be found in polynomial time. In terms of the economic properties, although the PDP is known to generate Pareto optimal outcomes, we show that, in some cases, which we identify, the SP is better for one of the two agents and also increases the social welfare.

Categories and Subject Descriptors

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

General Terms

Algorithms, Economics, Theory

Keywords

Game-theory, Negotiation, Approximation.

1. INTRODUCTION

Negotiation is a key form of interaction in multiagent systems. It is a process in which disputing agents decide how to divide the

Cite as: An Analysis of Feasible Solutions for Multi-Issue Negotiation Involving Nonlinear Utility Functions, Shaheen Fatima, Michael Wooldridge, Nicholas R. Jennings, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1041–1048

Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org), All rights reserved.

gains from cooperation [9]. This decision-making depends on factors such as the number of negotiators, the number of issues to be negotiated, the type of issues (i.e., divisible or indivisible), and the agents' utility functions [9, 13]. This paper focuses on bilateral multi-issue negotiation where issues are *divisible*, the agents have time constraints in the form of *deadlines* and *discount factors*, and the agents have different preferences over issues.

Given these differing preferences, Pareto-optimal agreements can be reached by negotiating all the issues together using a *package deal procedure* (PDP) [2]. In this procedure, the agents make alternating offers to each other and each such offer specifies a division on all the issues. The Pareto-optimality stems from the fact that the agents can make trade-offs across the issues and so an agent is willing to accept a lower utility on a less important issue, because it can obtain a higher utility on more important ones.

However, computing the equilibrium for the PDP is not always easy. Specifically, as we will show, if the agents' utilities are *nonlinear*¹ then the equilibrium strategies are hard to compute. This complexity represents a potentially significant barrier to the practical use of software negotiators. To overcome this problem, we explore two different solutions. The first is to approximate the nonlinear functions with linear ones and then use the PDP. The equilibrium for linear utilities can be computed in polynomial time. The other solution is to negotiate the issues in parallel but independently of each other using a *simultaneous procedure* (SP). Because the issues are dealt with independently, the equilibrium for SP can also be computed in polynomial time. However, since the two procedures differ in terms of their equilibrium solutions, our objective is to compare them in terms of their "computational properties" (i.e., their time complexities and approximation errors) and also their "economic properties" (i.e., the agents' utilities and the social welfare). In so doing, we show that an ϵ -Nash equilibrium for the PDP and the exact equilibrium for the SP can be computed in polynomial time. In terms of economic properties, although the PDP is known to generate Pareto optimal outcomes, we show that, in some cases, which we identify, the SP may be better for one of the two agents and also improve the social welfare.

Overall, this paper makes the following key contributions. First, it obtains, for the first time, ϵ -Nash equilibrium strategies for PDP and SP. Second, since the SP results in multiple equilibria, it formulates rules for reaching an outcome that maximizes social welfare. Third, it compares PDP and SP both in terms of their computational and their economic properties. When taken together, these results enable the designers of complex negotiation settings to make an informed choice about which procedure to use when.

¹Nonlinear utilities are common in many economic scenarios. For example, utility from money is nonlinear.

The remainder of the paper is organised as follows. Section 2 describes the negotiation setting. Section 3 obtains the equilibrium strategies for the PDP and analyzes its computational complexity. Section 4 introduces two solutions for overcoming the complexity of negotiating multiple issues with nonlinear utility functions. Section 5 provides a comparative analysis of these two solutions. Section 6 discusses related literature and Section 7 concludes.

2. THE NEGOTIATION SETTING

We extend the single issue model of [15, 2] to multiple issues. Before doing so, we give a brief overview of this model in terms of its equilibrium strategies. Then, we provide our analysis of this equilibrium to show how the negotiation deadline and discount factor affect it (the results of this analysis will subsequently be used to maximize the social welfare for the SP).

2.1 Dealing with a Single Issue

The single issue model of [2] (a variant of the one presented in [15]) is as follows. Two agents, a and b negotiate over a single divisible issue which is a ‘pie’ of size 1. The agents want to determine how to split it between themselves. Let $n \in \mathbb{N}^+$ be the deadline and $0 < \delta \leq 1$ the discount factor for both agents. The agents use an alternating offers protocol [12], which proceeds through a series of time periods. One of the agents, say a , starts in the first time period (i.e., $t = 1$) by making an offer ($0 \leq x \leq 1$) to b . Agent b can accept or reject the offer. If it accepts, negotiation ends in an agreement with a getting x and b getting $y = 1 - x$. Otherwise, negotiation proceeds to the next time period, in which agent b makes a counter-offer. This process continues until one of the agents either accepts an offer or quits negotiation (resulting in a conflict).

Let the pair $[x, y]$ denote the offer made at time t where x (y) denotes a 's (b 's) share. Then, the set of possible offers is $\{[x, y] | 0 \leq x \leq 1, \text{ and } x + y = 1\}$. The utility functions for a and b are:

$$UA(x, t) = \begin{cases} x \times \delta^{t-1} & \text{if } t \leq n \\ 0 & \text{otherwise} \end{cases}$$

$$UB(y, t) = \begin{cases} y \times \delta^{t-1} & \text{if } t \leq n \\ 0 & \text{otherwise} \end{cases}$$

The conflict utility (i.e., the utility received in the event that no deal is struck) is zero for both agents.

For this setting, the equilibrium offers were obtained in [2] as follows. Let $n = 1$ and let a be the first mover. If b accepts a 's proposal at $t = 1$, the division occurs as agreed; if not, neither agent gets anything (since the deadline is $n = 1$). Here, a is in a powerful position and is able to keep 100% of the pie and give nothing to b . Agent b accepts and an agreement takes place at $t = 1$.

Now, let $n = 2$ and $\delta = 1/2$. The first mover (say a) decides what to offer at $t = 1$, by looking ahead to $t = 2$ and reasoning backwards. Agent a reasons that if negotiation proceeds to $t = 2$, b will take 100 percent of the shrunken pie by offering $[0, 1/2]$. Thus, in the first time period, if a offers b anything less than $1/2$, b will reject the offer. So, at $t = 1$, agent a offers $[1/2, 1/2]$. Agent b accepts and an agreement occurs at $t = 1$.

In general, let $SA(t)$ ($SB(t)$) denote a 's (b 's) equilibrium strategy for time period t . Also, let $A_a(t)$ ($A_b(t)$) denote a 's (b 's) share in a 's equilibrium offer for t . The shares $B_a(t)$ ($B_b(t)$) in b 's equilibrium offer are defined analogously. Then the following strategies

²It is possible that b may reject such a proposal. However, irrespective of whether b accepts or rejects the proposal, it gets zero utility (because the deadline is $n = 1$). Thus, b accepts a 's offer.

form a Nash equilibrium [2]:

$$SA(t) = \begin{cases} \text{IF } a\text{'s TURN TO OFFER:} \\ \quad \text{OFFER } [\delta^{t-1} - B_b(t+1), B_b(t+1)] \\ \text{IF } a \text{ RECEIVES } [x, y]: \\ \quad \text{If } (UA(x, t) \geq U_t^a) \text{ ACCEPT} \\ \quad \text{else REJECT} \end{cases}$$

$$SB(t) = \begin{cases} \text{IF } b\text{'s TURN TO OFFER:} \\ \quad \text{OFFER } [\delta^{t-1} - A_a(t+1), A_a(t+1)] \\ \text{IF } b \text{ RECEIVES } [x, y]: \\ \quad \text{If } (UB(y, t) \geq U_t^b) \text{ ACCEPT} \\ \quad \text{else REJECT} \end{cases}$$

where $U_t^a = UA(A_a(t+1), t+1)$ and $U_t^b = UB(B_b(t+1), t+1)$. For the last time period $t = n$, we have $A_a(n) = \delta^{n-1}$, $A_b(n) = 0$, $B_a(n) = 0$, and $B_b(n) = \delta^{n-1}$. Given this, it is easy to verify using backward induction, that if a is the offering agent at $t < n$, we have:

$$A_a(t) = \sum_{i=0, j=t}^{j=n} (-1)^i \delta^{j-1} \quad ; \quad A_b(t) = \delta^{t-1} - A_a(t) \quad (1)$$

and if b is the offering agent at $t < n$, we have:

$$B_a(t) = \delta^{t-1} - B_b(t) \quad ; \quad B_b(t) = \sum_{i=0, j=t}^{j=n} (-1)^i \delta^{j-1} \quad (2)$$

The time to compute the equilibrium offer for $t = 1$ is $\mathcal{O}(n)$ [2].

Given the above equilibrium strategies, we now analyze how the deadline and the discount factor effect an agent's utility. Later, we will use the result of this analysis to find a way of maximizing social welfare for the SP.

Effect of deadline and discount factor on the equilibrium. The strategies $SA(t)$ and $SB(t)$ depend on two factors: the deadline (n) and the discount factor (δ). In order to analyze the effect of these factors on the negotiation outcome we vary δ between zero and one, and for each δ , determine the outcome for a range of n . The results of this analysis are shown in Figure 1. Here the continuous lines (dotted lines) denote the share for the first mover (second mover). In this figure, the plot for $\delta = 1$ shows the outcomes for $n \leq 10$. This is done to enhance clarity; the pattern remains unchanged for $n > 10$. The key results of this analysis are:

R_1 : For $\delta < 0.5$ and $n \geq 1$, the first mover's share is greater than that for the second, (i.e., if a is the first mover, then $A_a(1) > A_b(1)$, and if b is the first mover, then $B_b(1) > B_a(1)$).

R_2 : For $0.5 \leq \delta < 1$, there is a deadline $n_\delta \geq 1$ such that for all $n > n_\delta$, the first mover's share is strictly greater than that for the second mover.

For instance, for $\delta = 0.1$, $n_\delta = 1$, and for $\delta = 0.9$, $n_\delta = 40$ (see Figure 1). Later, in Theorem 4, we will use R_1 and R_2 to maximize social welfare for the SP.

We now extend this single issue model to multiple issues.

2.2 Dealing with Multiple Issues

Here a and b negotiate over $m > 1$ pies that represent m issues. Each pie is of size 1. The agents want to determine how to split each of the m pies. The discount factor is δ and the deadline is n for all the m issues. Let L_a (L_b) be an m -dimensional vector that represent a 's (b 's) share for the m issues. We assume that the utility




Figure 1: The effect of deadline (n) and discount factor (δ) on an agent's equilibrium share for single issue negotiation. The continuous lines denote the share for the first mover and the dotted lines those for the other agent.

functions are *separable*³ and nonlinear, i.e., UA and UB are of the following form:

$$UA(L_a, t) = \delta^{t-1} \sum_{c=1}^m FA_c(L_a^c) \quad (3)$$

$$UB(L_b, t) = \delta^{t-1} \sum_{c=1}^m FB_c(L_b^c) \quad (4)$$

where the functions FA_c and FB_c are nonlinear polynomials.

3. PACKAGE DEAL — EQUILIBRIUM

For the setting described in Section 2, we now analyze the strategic behavior of agents to determine equilibrium strategies for an incomplete information setting with nonlinear utilities. Before doing so, we give a brief overview of the equilibrium for complete information with linear utilities presented in [2].

3.1 The complete information setting

For the complete information setting, [2] obtained equilibrium for the case of linear utilities using backward induction as follows. Let $SAP(t)$ ($SBP(t)$) denote a 's (b 's) strategy for time t . For $t = n$, the offering agent gets a 100% of all the m shrunken pies. For all previous time periods, $t < n$, the offering agent (say a) proposes a package ($[L_a, L_b]$) such that b 's cumulative utility from it is equal to what b would get from its own offer for $t + 1$. If there is more than one package that gives b this utility, then a must find the one that maximizes its own cumulative utility. Thus, a must solve the following maximization (or trade-off) problem:

$$\begin{aligned} \text{TA :} \quad & \text{maximize} \quad UA(L_a, t) \quad (\text{or } \sum_{c=1}^m FA_c(L_a^c)) \\ & \text{subject to} \quad UA(L_b, t) = U_t^b \\ & \quad 0 \leq L_b^c \leq 1 \quad \text{and} \quad L_a^c = 1 - L_b^c \quad \text{for } 1 \leq c \leq m \end{aligned}$$

where U_t^b is b 's cumulative utility from its own offer $SBP(t+1)$. On the other hand, if a receives an offer $[L_a, L_b]$ at time t , then it accepts if $UA(L_a, t) = U_t^a$ where U_t^a is a 's cumulative utility from its own offer $SAP(t+1)$. The equilibrium strategy for b (in terms of TB) is defined analogously. Thus we have:

$$SAP(n) = \begin{cases} \text{OFFER } [\mathbf{1}, \mathbf{0}] & \text{IF } a\text{'s TURN TO OFFER} \\ \text{ACCEPT} & \text{IF } b\text{'s TURN TO OFFER} \end{cases}$$

³A function is said to be separable if it can be represented as the sum of several functions (generally nonlinear) of a single variable each [11]. Future work will deal with non-separable functions.

$$SBP(n) = \begin{cases} \text{OFFER } [\mathbf{0}, \mathbf{1}] & \text{IF } b\text{'s TURN TO OFFER} \\ \text{ACCEPT} & \text{IF } a\text{'s TURN TO OFFER} \end{cases}$$

where $\mathbf{0}$ ($\mathbf{1}$) denotes a vector of m zeros (ones). For all preceding time periods $t < n$, the strategies are as follows:

$$SAP(t) = \begin{cases} \text{OFFER } TA & \text{IF } a\text{'s TURN TO OFFER} \\ \text{IF } (UA(L_a, t) \geq U_t^a) \text{ ACCEPT} & \text{IF } a \text{ RECEIVES } [L_a, L_b] \\ \text{else REJECT} & \end{cases}$$

$$SBP(t) = \begin{cases} \text{OFFER } TB & \text{IF } b\text{'s TURN TO OFFER} \\ \text{IF } (UB(L_b, t) \geq U_t^b) \text{ ACCEPT} & \text{IF } b \text{ RECEIVES } [L_a, L_b] \\ \text{else REJECT} & \end{cases}$$

Complexity of TA or TB. For linear utilities, (i.e., UA and UB are linear) TA and TB are solvable in time linear in m [2].

3.2 The incomplete information setting

One of the most common sources of uncertainty is with regard to utility functions. The agents are uncertain about their utilities and we model this uncertainty as follows. There are r possible cumulative utility functions for each agent. For agent a , these are denoted UA^1, \dots, UA^r , and for b , they are UB^1, \dots, UB^r . Let α_a denote a discrete probability distribution function over the UA^1, \dots, UA^r . and α_b that over UB^1, \dots, UB^r . The distribution functions α_a and α_b , and the utility functions UA^1, \dots, UA^r , and UB^1, \dots, UB^r are common knowledge to a and b .

Because of this uncertainty, agents can only compute their *expected utilities*. Let $EUA(L_a, t)$ and $EUB(L_b, t)$ denote the expected utility that a and b get from L_a and L_b at time t . These utilities are defined as follows:

$$EUA(L_a, t) = \delta^{t-1} \left(\sum_{d=1}^r \alpha_a(d) \times \sum_{c=1}^m FA_c^d(L_a^c) \right) \quad (5)$$

$$EUB(L_b, t) = \delta^{t-1} \left(\sum_{d=1}^r \alpha_b(d) \times \sum_{c=1}^m FB_c^d(L_b^c) \right) \quad (6)$$

As before, FA^d and FB^d are nonlinear. Given this, agent a 's trade-off problem at time t is to find a package $[L_a, L_b]$ that solves the following maximization problem:

$$\begin{aligned} \text{TA-I} \quad & \text{maximize} \quad \sum_{d=1}^r \alpha_a(d) \times \sum_{c=1}^m FA_c^d(L_a^c) \\ & \text{subject to} \quad \sum_{d=1}^r \alpha_b(d) \times \sum_{c=1}^m FB_c^d(L_b^c) = U_t^b \\ & \quad 0 \leq L_a^c \leq 1 \quad \text{for } 1 \leq c \leq m \quad \text{and} \quad L_a^c = 1 - L_b^c \end{aligned}$$

Here U_t^b denotes b 's expected utility from its own offer for time $t + 1$. If FA^d and FB^d are nonlinear, then EUA and EUB are nonlinear since EUA is a linear transformation of FA and EUB that of FB. For the nonlinear case, a 's equilibrium strategy for time t is the same as SAP(t) (defined in Section 3.1) with TA replaced with TA-I, and UA replaced with EUA. Likewise for agent b . TA-I and TB-I are both nonlinear optimization problems with both the objective function and constraint being nonlinear.

Complexity of TA-I and TB-I. For the nonlinear utilities (defined in Equations 3 and 4) TA-I and TB-I are optimization problems with nonlinear objectives and nonlinear constraints and such problems are, in general, computationally hard. This is because, a specific instance of TA-I (or TB-I), viz., the *quadratic optimization* problem (which has a quadratic objective function and a linear constraint) is NP-hard [5]. In more detail, for a quadratic instance of TA-I, FA_c^d is quadratic and FB_c^d is linear. Likewise for TB-I.

4. TOWARDS FEASIBILITY

In order to overcome the complexity of TA-I and TB-I, we explore two solutions: approximating nonlinear utilities with linear ones and using the PDP (Section 4.1), and using the SP (Section 4.2).

4.1 PDP with approximate linear utilities

The source of complexity in computing equilibrium strategies for the PDP is the complexity of solving the trade-off problem, which is a nonlinear optimization problem. One approach to solving such problems is to approximate the nonlinear functions involved — objective function and/or constraints — with straight line segments and then use methods that are appropriate for linear problems [11]. The resulting linear optimization problem can then be solved in polynomial time. In more detail, in a completely general nonlinear optimization problem, linear approximations are not easy to specify algebraically. However, with separable functions, a mathematical specification is more easily done [11]. As per Equations 3 and 4, the objective and constraint in TA-I or TB-I are separable. Given this, let \bar{U}_a , \bar{F}_a , and \bar{EUA} denote the linear approximations for the functions UA, FA, and EUA respectively. The approximation functions for agent b (i.e., \bar{U}_b , \bar{F}_b , and \bar{EUB}) are denoted analogously. We assume that \bar{F}_a^d (\bar{F}_b^d) is increasing in L_a^c (L_b^c). Being linear, these approximate functions are of the following form:

$$\bar{F}_a^d(L_a^c) = \mathcal{A}_c^d L_a^c \quad \text{and} \quad \bar{F}_b^d(L_b^c) = \mathcal{B}_c^d L_b^c \quad (7)$$

where $\mathcal{A}_c^d \in \mathbb{R}_+$ and $\mathcal{B}_c^d \in \mathbb{R}_+$ denote the "weights" that a and b associate with the different issues in their approximate utility functions.

Note that it is not our objective to find the approximate functions. Rather, given the functions \bar{F}_a^d , \bar{F}_b^d , and the associated approximation errors, our objective is to find an approximate equilibrium. To this end, let η denote an upper bound on the "absolute error" in the function \bar{F}_a^d or \bar{F}_b^d , i.e.,

$$abs(\mathcal{A}_c^d L_a^c - \bar{F}_a^d(L_a^c)) \leq \eta$$

and

$$abs(\mathcal{B}_c^d L_b^c - \bar{F}_b^d(L_b^c)) \leq \eta.$$

Given \bar{F}_a and \bar{F}_b , the approximate expected utilities are:

$$\bar{EUA}(L_a, t) = \delta^{t-1} \left(\sum_{d=1}^r \alpha_a(d) \times \sum_{c=1}^m \bar{F}_a^d(L_a^c) \right) \quad (8)$$

$$\bar{EUB}(L_b, t) = \delta^{t-1} \left(\sum_{d=1}^r \alpha_b(d) \times \sum_{c=1}^m \bar{F}_b^d(L_b^c) \right) \quad (9)$$

In order to obtain the error in \bar{EUA} or \bar{EUB} in terms of η , we use standard error propagation rules [16]. This propagation gives us $abs(\bar{EUA}(L_a, t) - \bar{EUA}(L_a, t)) \leq \delta^{t-1} \sum_{d=1}^r \alpha_a(d) \times m\eta = \delta^{t-1} m\eta$. Likewise, $abs(\bar{EUB}(L_b, t) - \bar{EUB}(L_b, t)) \leq \delta^{t-1} m\eta$.

Given the above linear approximations, we find an approximate equilibrium for the PDP. Let $\bar{SAP}(t)$ ($\bar{SBP}(t)$) denote a 's (b 's) approximate equilibrium strategy for time period t . Here agent a 's trade-off or optimization problem for time t is called $\bar{TA-I}$ and is:

$$\begin{aligned} \bar{TA-I} \quad & \text{maximize} \quad \sum_{d=1}^r \alpha_a(d) \times \sum_{c=1}^m \mathcal{A}_c^d L_a^c \\ & \text{subject to} \quad \sum_{d=1}^r \alpha_a(d) \times \sum_{c=1}^m \mathcal{B}_c^d L_a^c = \bar{U}_t^b \\ & 0 \leq L_a^c \leq 1 \quad \text{for } 1 \leq c \leq m \text{ and } L_a^c = 1 - L_b^c \end{aligned}$$

where \bar{U}_t^b denotes b 's approximate expected cumulative utility from $\bar{SBP}(t+1)$. Agent b 's optimization problem ($\bar{TB-I}$) is defined analogously in terms of \bar{U}_t^a . Note that both $\bar{TA-I}$ and $\bar{TB-I}$ are linear knapsack problems and therefore solvable in $\mathcal{O}(m)$ time [10].

Since $\bar{TA-I}$ will give an approximately optimal solution, the error in it must be measured with respect to the corresponding exactly optimal solution. The exact counterpart of the approximate optimization problem $\bar{TA-I}$ is called $\bar{ETA-I}$ and is:

$$\begin{aligned} \bar{ETA-I} \quad & \text{maximize} \quad \sum_{d=1}^r \alpha_a(d) \times \sum_{c=1}^m \mathcal{F}_a^d L_a^c \\ & \text{subject to} \quad \sum_{d=1}^r \alpha_a(d) \times \sum_{c=1}^m \mathcal{B}_c^d L_a^c = \bar{U}_t^b \\ & 0 \leq L_a^c \leq 1 \quad \text{for } 1 \leq c \leq m \text{ and } L_a^c = 1 - L_b^c \end{aligned}$$

For b , $\bar{ETB-I}$ is defined analogously. In order to find the difference between the solutions to $\bar{ETA-I}$ and $\bar{TA-I}$, consider time t and let a be the offering agent. Let $[O_a, O_b]$ (where O_a and O_b are m element vectors that contain a 's and b 's shares for the m issues) denote the solution to $\bar{ETA-I}$. Likewise, let $[\bar{O}_a, \bar{O}_b]$ denote the solution to $\bar{TA-I}$. Then, let $F = \bar{EUA}(O_a, t)$ and $\bar{F} = \bar{EUA}(\bar{O}_a, t)$, i.e., F is the exact maximum, \bar{F} is the exact value of the function \bar{EUA} at the approximately optimal solution \bar{O}_a . Finally, let ϵ be an upper bound on $abs(F - \bar{F})$. Likewise, for b , let ϵ denote an upper bound on $abs(\bar{EUB}(\bar{O}_b, t) - \bar{EUB}(O_b, t))$. Theorem 1 characterizes an ϵ -approximate Nash equilibrium (a strategy profile is said to form an ϵ -approximate Nash equilibrium if no player can gain more than ϵ by deviating) for the PDP.

THEOREM 1. *The following strategies form an ϵ -approximate Nash equilibrium for time t where $\epsilon = 2m\eta\delta^{t-1}$. For $t = n$, $\bar{SAP}(n) = \text{SAP}(n)$ and $\bar{SBP}(n) = \text{SBP}(n)$. For all preceding time periods $t < n$, the strategies are as follows:*

$$\bar{SAP}(t) = \begin{cases} \text{OFFER } \bar{TA-I} & \text{IF } a \text{'s TURN TO OFFER} \\ \text{If } (\bar{EUB}(L_b, t) \geq \bar{U}_t^b) \text{ ACCEPT} & \text{IF } a \text{ RECEIVES } [L_a, L_b] \\ \text{else REJECT} & \end{cases}$$

$$\bar{SBP}(t) = \begin{cases} \text{OFFER } \bar{TB-I} & \text{IF } b \text{'s TURN TO OFFER} \\ \text{If } (\bar{EUA}(L_a, t) \geq \bar{U}_t^a) \text{ ACCEPT} & \text{IF } b \text{ RECEIVES } [L_a, L_b] \\ \text{else REJECT} & \end{cases}$$

An agreement on all the issues occurs at $t = 1$.

PROOF. *It is straightforward to obtain $\bar{SAP}(t)$ and $\bar{SBP}(t)$ using backward induction. Thus we will now prove that $\epsilon = 2m\eta\delta^{t-1}$.*

Consider time period t and let a be the offering agent. Let the definitions for $[O_a, O_b]$, $[\bar{O}_a, \bar{O}_b]$, F , and \bar{F} be as given earlier. In

addition, let $F_1 = \overline{\text{EUA}}(O_a, t)$ and $F_2 = \overline{\text{EUA}}(\overline{O}_a, t)$. In other words, F_1 is the approximate value of EUA at the exact optimal solution O_a , and F_2 is the approximate value of EUA at the approximately optimal solution \overline{O}_a . We need to find an ϵ such that $F = \overline{F} \pm \epsilon$. To do so, assume that $O_a \neq \overline{O}_a$. This implies that $\overline{F} \leq F$, because F is the exact maximum and so all other values of the function EUA must be no greater than F . Recall that the upper bound on the approximation error in $\overline{\text{EUA}}$ is $m\eta\delta^{t-1}$. Thus, we have $F_2 - m\eta\delta^{t-1} \leq \overline{F} \leq F_2 + m\eta\delta^{t-1}$. Substituting this relation in $\overline{F} \leq F$, we get $F_2 - m\eta\delta^{t-1} \leq F$ or $F_2 \leq F + m\eta\delta^{t-1}$. Again, since $F_2 \geq \overline{F} - m\eta\delta^{t-1}$, we get $\overline{F} - m\eta\delta^{t-1} \leq F + m\eta\delta^{t-1}$ or $F \geq \overline{F} - 2m\eta\delta^{t-1}$. Thus, we have a lower bound on F .

We now need to find an upper bound on F . To do so, we first find the relation between F_1 and F_2 . We prove that $F_1 \leq F_2$ by contradiction. Assume that $(F_1 = \overline{\text{EUA}}(O_a, t)) > (\overline{\text{EUA}}(\overline{O}_a, t) = F_2)$. But $F_1 > F_2$ contradicts the fact that \overline{O}_a is the optimal solution to the linear knapsack problem $(\overline{\text{TA}}-1)$ and \overline{O}_a therefore maximizes $\overline{\text{EUA}}$. In other words, the relation $F_1 > F_2$ has led to a contradiction. We therefore have $F_1 \leq F_2$. It follows that $F_1 - m\eta\delta^{t-1} \leq F_2 - m\eta\delta^{t-1}$. We already know that $\overline{F} \geq F_2 - m\eta\delta^{t-1}$. This relation together with $F_1 - m\eta\delta^{t-1} \leq F_2 - m\eta\delta^{t-1}$ gives $F_1 \leq \overline{F} + m\eta\delta^{t-1}$. Recall that F_1 is the value of the approximate function $\overline{\text{EUA}}$ which has error $m\eta\delta^{t-1}$. So we have $F_1 - m\eta\delta^{t-1} \leq F \leq F_1 + m\eta\delta^{t-1}$ or $F_1 \geq F - m\eta\delta^{t-1}$. The relations $F_1 \leq \overline{F} + m\eta\delta^{t-1}$ and $F_1 \geq F - m\eta\delta^{t-1}$ together give $F \leq \overline{F} + 2m\eta\delta^{t-1}$. Thus, $\overline{F} - 2m\eta\delta^{t-1} \leq F \leq \overline{F} + 2m\eta\delta^{t-1}$. We now have both an upper and a lower bound on F . In other words, $\epsilon = 2m\eta\delta^{t-1}$ for the case when a is the offering agent.

In the same way, it can be shown that $\epsilon = 2m\eta\delta^{t-1}$ for the case where b is the offering agent. \square

Finally, as mentioned earlier, $\overline{\text{TA}}-1$ or $\overline{\text{TB}}-1$ can be solved in $\mathcal{O}(m)$ time, so the time to compute $\overline{\text{SAP}}(1)$ or $\overline{\text{SBP}}(1)$ is $\mathcal{O}(mn)$.

4.2 Simultaneous procedure

Another way of overcoming the complexity of nonlinear optimization is to use the *simultaneous* procedure which works as follows. For each issue, one of the negotiators is chosen as the first mover. Since the first mover can be different for different issues, there are 2^m possible ways from which to choose a first mover for the m issues. Once this choice is made, negotiation on each of the m issues begins in the first time period and each issue is negotiated independently of all other issues. So there are m single issue negotiations held in parallel. The deadline for each of these negotiations is n .

As the issues are dealt with independently, it is optimal for an agent to maximize its utility for each individual issue. Let $\text{SAS}_c(t)$ ($\text{SBS}_c(t)$) denote a 's (b 's) equilibrium strategy for issue c for time t . Thus, at t , a must solve the following optimization problem:

$$\begin{aligned} \text{TAS-I}_c \quad & \text{maximize} \quad \sum_{d=1}^r \alpha_a(d) \times \text{FA}_c^d L_a^c \\ & \text{subject to} \quad \sum_{d=1}^r \alpha_a(d) \times \text{FB}_c^d(L_a^c) = U_t^b \\ & 0 \leq L_a^c \leq 1 \quad \text{for } 1 \leq c \leq m \text{ and } L_a^c = 1 - L_b^c \end{aligned}$$

where U_t^b is b 's utility from its offer $\text{SBS}_c(t+1)$. For b , TAS-I_c is defined analogously. Note that $U_{n+1}^a = U_{n+1}^b = 0$.

Theorem 2 formulates the equilibrium strategies for the SP.

THEOREM 2. *The following strategies form a Nash equilibrium*

for issue c ($1 \leq c \leq m$) for time period t ($1 \leq t \leq n$):

$$\begin{aligned} \text{SAS}_c(t) &= \begin{cases} \text{OFFER TAS-I}_c & a \text{'s TURN TO OFFER} \\ \text{If } \text{EUA}_c(L_a^c, t) \geq U_t^a \text{ ACCEPT} & a \text{ RECEIVES } [L_a^c, L_b^c] \\ \text{else REJECT} \end{cases} \\ \text{SBS}_c(t) &= \begin{cases} \text{OFFER TBS-I}_c & b \text{'s TURN TO OFFER} \\ \text{If } \text{EUB}_c(L_b^c, t) \geq U_t^b \text{ ACCEPT} & b \text{ RECEIVES } [L_a^c, L_b^c] \\ \text{else REJECT} \end{cases} \end{aligned}$$

An agreement on all the issues takes place in the first time period.

PROOF. As for linear utilities (see Section 2.1). \square

Note that, solving TAS-I_c or TBS-I_c takes $\mathcal{O}(1)$ time. Since TAS-I_c or TBS-I_c must be solved for n time periods for each of the m issues, the time to compute an offer for $t = 1$ is $\mathcal{O}(mn)$.

Recall that the time to compute an equilibrium offer for $t = 1$ for the PDP is $\mathcal{O}(mn)$. Thus, the PDP for approximate linear utilities and the SP for the given nonlinear utilities are identical in terms of the time required to compute an equilibrium offer. However, they differ in terms of their equilibrium outcomes i.e., the individual agents' utilities, the Pareto optimality, and the social welfare of the outcome. Since the PDP (unlike the SP) is known to generate Pareto optimal outcomes, we compare the two procedures in terms of the agents' utilities and the social welfare. Clearly, this comparison depends on the agents' utility functions. If we compare the procedures for one specific utility function, we cannot draw any general conclusions regarding the performance of the two procedures. Thus, instead of considering one specific nonlinear utility function, we do the following. We compare the two procedures in terms of the approximate linear utilities. The weights for these linear utilities are generated randomly as per a given probability distribution function (details in Section 5). This will allow us to do the comparison for a wide range of utility functions and thereby draw some general conclusions about the two procedures. To this end, we first formulate approximate equilibrium strategies for the SP. Let $\overline{\text{SAS}}_c(t)$ ($\overline{\text{SBS}}_c(t)$) denote a 's (b 's) approximate equilibrium strategy for issue c for time period t .

For a 's optimization problem that corresponds to TAS-I_c is denoted $\overline{\text{TAS-I}}_c$ and is of the following form:

$$\begin{aligned} \overline{\text{TAS-I}}_c \quad & \text{maximize} \quad \sum_{d=1}^r \alpha_a(d) \times \overline{\text{FA}}_c^d L_a^c \quad \text{or} \quad L_a^c \\ & \text{subject to} \quad \sum_{d=1}^r \alpha_a(d) \times \overline{\text{FB}}_c^d(L_a^c) = \overline{U}_t^b \\ & 0 \leq L_a^c \leq 1 \quad \text{for } 1 \leq c \leq m \text{ and } L_a^c = 1 - L_b^c \end{aligned}$$

where \overline{U}_t^b is b 's approximate utility from its own offer $\overline{\text{SBS}}_c(t+1)$. For b , $\overline{\text{TAS-I}}_c$ is defined analogously in terms of \overline{U}_t^a . Theorem 3 gives approximate equilibrium strategies for the SP.

THEOREM 3. *The following strategies form an ϵ -approximate Nash equilibrium for time period t where $\epsilon = 2m\eta\delta^{t-1}$. An agreement on all the issues occurs at $t = 1$. Here, $\overline{\text{SAS}}_c(t) = \text{SA}(t)$ and $\overline{\text{SBS}}_c(t) = \text{SB}(t)$ for issue c .*

PROOF. As per Equation 7, $\overline{\text{FB}}_c^d$ ($\overline{\text{FB}}_c^d$) is increasing in L_a^c (L_b^c). Also, since the issues are negotiated independently, the optimal strategy for each agent is to maximize its share for each individual. We therefore get $\overline{\text{SAS}}_c(t) = \text{SA}(t)$, and $\overline{\text{SBS}}_c(t) = \text{SB}(t)$ for issue c (see Section 2.1 for a definition of $\text{SA}(t)$ and $\text{SB}(t)$ in terms of A_a , A_b , B_a , and B_b). We will now find ϵ (an upper bound on the error in an agent's approximate cumulative utility).

For time t , Theorem 1 gives $\epsilon = 2m\eta\delta^{t-1}$ as the upper bound for the error in an agent's cumulative utility. Thus, for $m = 1$ we get $\epsilon = 2\eta\delta^{t-1}$. Since there are $m > 1$ issues, the upper bound on the error in an agent's cumulative utility is $\epsilon = 2m\eta\delta^{t-1}$. \square

Note that the outcome for the SP (for both, the exact equilibrium of Theorem 2 and the approximate equilibrium of Theorem 3) depends on the choice of the first mover for each of the m issues. This choice affects the agent's utilities and also the social welfare. Given this, it is important to find the maximum possible social welfare and show how the first mover for each issue must be chosen in order to achieve this maximum.

Maximum Possible Social Welfare: Since there are multiple outcomes, we determine the maximum possible approximate social welfare (measured in terms of the sum of the agents' approximate cumulative utilities) that can be obtained in equilibrium. Let SW_U denote this maximum. Recall from Section 2 (results R_1 and R_2), that, in most cases, the first mover gets a greater share of the pie. In other words, the outcome for the SP depends on the choice of the first mover for each of the m issues.

We first introduce some notation and then obtain SW_U . Let L and \bar{L} each denote m element vectors. Also, for a given c , let $\beta = \bar{L}^c - L^c$ and $\Delta\overline{EUA}_\beta^c$ and $\Delta\overline{EUB}_\beta^c$ be defined as follows:

$$\Delta\overline{EUA}_\beta^c = \overline{EUA}(\bar{L}^c, t) - \overline{EUA}(L^c, t) \quad (10)$$

$$\Delta\overline{EUB}_\beta^c = \overline{EUB}(\bar{L}^c, t) - \overline{EUB}(L^c, t) \quad (11)$$

Then, a is said to 'value' issue c more than b if, for some $\beta > 0$, $\Delta\overline{EUA}_\beta^c > \Delta\overline{EUA}_\beta^b$. Theorem 4 uses $\Delta\overline{EUA}_\beta^c$ and $\Delta\overline{EUA}_\beta^b$ to formulate rules (\mathcal{R}_1 and \mathcal{R}_2) to maximize the approximate welfare.

THEOREM 4. For the SP, let \mathcal{R}_1 and \mathcal{R}_2 be two rules for choosing the first mover for issue c ($1 \leq c \leq m$) where:

\mathcal{R}_1 : If $\Delta\overline{EUA}_\beta^c > \Delta\overline{EUB}_\beta^c$ then a is the first mover for issue c , else if $\Delta\overline{EUA}_\beta^c < \Delta\overline{EUB}_\beta^c$ b is the first mover.

\mathcal{R}_2 : If $\Delta\overline{EUA}_\beta^c = \Delta\overline{EUB}_\beta^c$ then the first mover is chosen randomly.

If ($\delta < 0.5$) or ($\delta > 0.5$ and $n_c < n_\delta$), then the social welfare for the ϵ -Nash equilibrium outcome for rules \mathcal{R}_1 and \mathcal{R}_2 is SW_U .

PROOF. Since the issues are dealt with independently, the welfare from all the m issues is maximized by maximizing the welfare for each individual issue. From Theorem 3 and Figure 1, we know that, for a given n and δ , the approximate equilibrium outcome for an issue depends on the first mover. Consider issue c . There are two possible outcomes for this issue: one that corresponds to a being the first mover and the other to b being the first mover. For the former, a 's and b 's shares are $A_a(1)$ and $A_b(1)$ respectively, while for the latter, they are $B_b(1)$ and $B_a(1)$. From Figure 1, we know that $A_a(1) > A_b(1)$ and $B_b(1) > B_a(1)$. In order to show that \mathcal{R}_1 and \mathcal{R}_2 give social welfare SW_U , we let $\bar{L}^c = A_a(1)$, $L^c = A_b(1)$. As per Equations 1 and 2, we have $A_a(1) = B_b(1)$ and $A_b(1) = B_a(1)$. This implies $\bar{L}^c = A_a(1) = B_b(1)$ and $L^c = A_b(1) = B_a(1)$. Then, let $\beta = \bar{L}^c - L^c$ and consider the following three possible relations between $\Delta\overline{EUA}_\beta^c$ and $\Delta\overline{EUB}_\beta^c$.

1. $\Delta\overline{EUA}_\beta^c > \Delta\overline{EUB}_\beta^c$: From Equations 10 and 11, we have the following relation:

$$\overline{EUA}(\bar{L}^c, 1) - \overline{EUA}(L^c, 1) > \overline{EUB}(\bar{L}^c, 1) - \overline{EUB}(L^c, 1)$$

This can be rewritten as:

$$\overline{EUA}(\bar{L}^c, 1) + \overline{EUB}(L^c, 1) > \overline{EUB}(\bar{L}^c, 1) + \overline{EUA}(L^c, 1).$$

Thus, the social welfare if a is the first mover (i.e., $\overline{EUA}(\bar{L}^c, 1) + \overline{EUB}(L^c, 1)$) is higher than the welfare if b is the first mover (i.e., $\overline{EUB}(\bar{L}^c, 1) + \overline{EUA}(L^c, 1)$), so $SW_U^c = \overline{EUA}(\bar{L}^c, 1) + \overline{EUB}(L^c, 1)$ and the rule \mathcal{R}_1 gives this welfare.

2. $\Delta\overline{EUA}_\beta^c = \Delta\overline{EUB}_\beta^c$: From Equations 10 and 11, we have the following relation:

$$\overline{EUA}(\bar{L}^c, 1) - \overline{EUA}(L^c, 1) = \overline{EUB}(\bar{L}^c, 1) - \overline{EUB}(L^c, 1)$$

This can be rewritten as:

$$\overline{EUA}(\bar{L}^c, 1) + \overline{EUB}(L^c, 1) = \overline{EUB}(\bar{L}^c, 1) + \overline{EUA}(L^c, 1)$$

Thus, the social welfare if a is the first mover is equal to the social welfare if b is the first mover, so $SW_U^c = \overline{EUA}(\bar{L}^c, 1) + \overline{EUB}(L^c, 1) = \overline{EUB}(\bar{L}^c, 1) + \overline{EUA}(L^c, 1)$ and \mathcal{R}_2 gives SW_U^c .

3. $\Delta\overline{EUA}_\beta^c < \Delta\overline{EUB}_\beta^c$: From Equations 10 and 11, we have one of the following relations:

$$\overline{EUA}(\bar{L}^c, 1) - \overline{EUA}(L^c, 1) < \overline{EUB}(\bar{L}^c, 1) - \overline{EUB}(L^c, 1)$$

or

$$\overline{EUA}(\bar{L}^c, 1) + \overline{EUB}(L^c, 1) < \overline{EUB}(\bar{L}^c, 1) + \overline{EUA}(L^c, 1)$$

In other words, the social welfare if b is the first mover is higher than the social welfare if a is the first mover and the rule \mathcal{R}_1 chooses b as the first mover.

Thus, \mathcal{R}_1 and \mathcal{R}_2 always give an outcome with welfare SW_U . \square

5. COMPARISON BETWEEN PDP AND SP

In order to compare the PDP and the SP in terms of the utilities they yield to the agents and also in terms of their social welfare, we conducted the following experimental analysis. Since the outcomes for the two procedures depend on m , n , δ , and the players' weights (A and B defined in Equation 7) for the different issues, we consider a wide range of settings by varying these four parameters. For these settings, we compare the PDP and the SP as follows.

Let a be the first mover for the PDP. Then, unlike the PDP, the first mover for the SP may be the same or different for different issues. Specifically, for the SP, we consider the following two cases:

C1: Agent a is the first mover for each issue

C2: The first mover for each issue is chosen as per the rules \mathcal{R}_1 and \mathcal{R}_2

Note that, for both cases, a is the first mover for the PDP. Since there is only one first mover for the PDP, case C1 compares the PDP with the SP by having a as the first mover for PDP and also for each issue for the SP. In contrast, case C2 compares the best outcome (best in terms of social welfare) for the SP with the outcome for PDP⁴.

At this stage, we would like to point out that we will compare the two procedures in terms of their approximate equilibria. Recall from Section 4, that ϵ is the upper bound on the error in the approximate equilibrium utilities. Hence, in what follows, we will quantify the approximation error in the results of our comparative analysis in terms of ϵ .

⁴For our experimental analysis we also considered the case where b is the first mover for the PDP and for each issue for the SP. Since the players are symmetric, the results for this case are the same as those for the case where a is the first mover for all the issues. Hence, we only describe the case with a as the first mover.

δ	$m = 50$			$m = 65$			$m = 75$			$m = 85$			$m = 100$		
	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}
0.1	7.98 $\pm\epsilon$	8.1 $\pm\epsilon$	7.99 $\pm 1.8\epsilon$	8.05 $\pm 0.7\epsilon$	8 $\pm 0.7\epsilon$	8 $\pm 1.4\epsilon$	8.31 $\pm 0.6\epsilon$	8.18 $\pm 0.6\epsilon$	8.3 $\pm 1.2\epsilon$	8.14 $\pm 0.5\epsilon$	8.2 $\pm 0.5\epsilon$	8.14 $\pm\epsilon$	8.35 $\pm 0.5\epsilon$	8.14 $\pm 0.5\epsilon$	8.33 $\pm 0.9\epsilon$
0.3	17.8 $\pm 0.5\epsilon$	17.8 $\pm 0.5\epsilon$	17.8 $\pm 0.7\epsilon$	17.8 $\pm 0.4\epsilon$	17.65 $\pm 0.4\epsilon$	17.8 $\pm 0.6\epsilon$	17.7 $\pm 0.3\epsilon$	17.8 $\pm 0.3\epsilon$	17.7 $\pm 0.5\epsilon$	17.9 $\pm 0.3\epsilon$	17.7 $\pm 0.3\epsilon$	17.8 $\pm 0.4\epsilon$	18.13 $\pm 0.3\epsilon$	17.8 $\pm 0.3\epsilon$	18 $\pm 0.4\epsilon$
0.5	21.9 $\pm 0.4\epsilon$	22 $\pm 0.4\epsilon$	21.99 $\pm 0.6\epsilon$	24.4 $\pm 0.3\epsilon$	22.2 $\pm 0.3\epsilon$	22.35 $\pm 0.4\epsilon$	22.9 $\pm 0.3\epsilon$	21.25 $\pm 0.3\epsilon$	22.36 $\pm 0.4\epsilon$	23 $\pm 0.2\epsilon$	22.9 $\pm 0.2\epsilon$	23 $\pm 0.3\epsilon$	22.57 $\pm 0.2\epsilon$	22.89 $\pm 0.2\epsilon$	22.69 $\pm 0.3\epsilon$
0.7	23.7 $\pm 0.49\epsilon$	23.3 $\pm 0.49\epsilon$	23.56 $\pm 0.54\epsilon$	24.9 $\pm 0.3\epsilon$	22.18 $\pm 0.3\epsilon$	23.78 $\pm 0.4\epsilon$	26.5 $\pm 0.3\epsilon$	21.8 $\pm 0.3\epsilon$	24.59 $\pm 0.4\epsilon$	26.5 $\pm 0.2\epsilon$	21.1 $\pm 0.2\epsilon$	24.3 $\pm 0.3\epsilon$	24.77 $\pm 0.2\epsilon$	21.2 $\pm 0.2\epsilon$	23.26 $\pm 0.3\epsilon$
0.9	26.18 $\pm 0.5\epsilon$	22.83 $\pm 0.5\epsilon$	24.62 $\pm 0.6\epsilon$	28 $\pm 0.3\epsilon$	21.88 $\pm 0.3\epsilon$	25.12 $\pm 0.4\epsilon$	28.08 $\pm 0.3\epsilon$	20.95 $\pm 0.3\epsilon$	24.64 $\pm 0.4\epsilon$	28.12 $\pm 0.2\epsilon$	21.27 $\pm 0.2\epsilon$	24.87 $\pm 0.3\epsilon$	25.63 $\pm 0.2\epsilon$	24.6 $\pm 0.2\epsilon$	25.17 $\pm 0.25\epsilon$

Table 1: A comparison of the two procedures for uniformly distributed weights for the case C1.

δ	$m = 50$			$m = 65$			$m = 75$			$m = 85$			$m = 100$		
	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}	Δ_{UA}	Δ_{UB}	Δ_{SW}
0.1	33.3 $\pm 0.3\epsilon$	-503 $\pm 0.3\epsilon$	-14.1 $\pm 1\epsilon$	36.8 $\pm 0.2\epsilon$	-535 $\pm 0.2\epsilon$	-16.4 $\pm 0.7\epsilon$	35.84 $\pm 0.2\epsilon$	-532 $\pm 0.2\epsilon$	-16.5 $\pm 0.6\epsilon$	34.55 $\pm 0.2\epsilon$	-519 $\pm 0.2\epsilon$	-14.8 $\pm 0.6\epsilon$	38.6 $\pm 0.1\epsilon$	-549 $\pm 0.1\epsilon$	-15.3 $\pm 0.5\epsilon$
0.3	37.95 $\pm 0.3\epsilon$	-113 $\pm 0.3\epsilon$	3.7 $\pm 4\epsilon$	36.8 $\pm 0.2\epsilon$	-106 $\pm 0.2\epsilon$	3.9 $\pm 4\epsilon$	36 $\pm 0.2\epsilon$	-105 $\pm 0.2\epsilon$	4.46 $\pm 0.6\epsilon$	36.84 $\pm 0.2\epsilon$	-108 $\pm 0.2\epsilon$	4.2 $\pm 4\epsilon$	37.7 $\pm 0.1\epsilon$	-111 $\pm 0.1\epsilon$	3.14 $\pm 2\epsilon$
0.5	35.2 $\pm 0.3\epsilon$	-31 $\pm 0.3\epsilon$	12.5 $\pm\epsilon$	34.7 $\pm 0.2\epsilon$	-28 $\pm 0.2\epsilon$	13.9 $\pm 0.7\epsilon$	35.6 $\pm 0.2\epsilon$	-30 $\pm 0.2\epsilon$	13.43 $\pm 0.6\epsilon$	35.3 $\pm 0.2\epsilon$	-28 $\pm 0.2\epsilon$	14.4 $\pm 0.5\epsilon$	35.6 $\pm 0.2\epsilon$	-29 $\pm 0.2\epsilon$	13.58 $\pm 0.5\epsilon$
0.7	33.16 $\pm 0.4\epsilon$	3.85 $\pm 0.4\epsilon$	20.6 $\pm 0.6\epsilon$	33.1 $\pm 0.3\epsilon$	0.01 $\pm 0.3\epsilon$	19 $\pm 0.5\epsilon$	32.6 $\pm 0.2\epsilon$	4.4 $\pm 0.2\epsilon$	20.9 $\pm 0.4\epsilon$	31.2 $\pm 0.2\epsilon$	4.6 $\pm 0.2\epsilon$	20.55 $\pm 0.4\epsilon$	32.4 $\pm 0.2\epsilon$	3.7 $\pm 0.2\epsilon$	20.5 $\pm 0.3\epsilon$
0.9	27.8 $\pm 0.4\epsilon$	21.8 $\pm 0.4\epsilon$	24.95 $\pm 0.5\epsilon$	26.9 $\pm 0.3\epsilon$	18.8 $\pm 0.3\epsilon$	23.03 $\pm 0.4\epsilon$	28.4 $\pm 0.3\epsilon$	20.1 $\pm 0.3\epsilon$	24.4 $\pm 0.4\epsilon$	21.8 $\pm 0.2\epsilon$	17.6 $\pm 0.2\epsilon$	23.2 $\pm 0.3\epsilon$	27.6 $\pm 0.2\epsilon$	18.11 $\pm 0.2\epsilon$	23.2 $\pm 0.3\epsilon$

Table 2: A comparison of the two procedures for uniformly distributed weights for the case C2.

We first introduce notation and then describe the experimental setting. For $t = 1$, let the approximate equilibrium offer for the PDP be $[L_a, L_b]$, for the SP for case C1 let it be $[\bar{L}_a, \bar{L}_b]$, and for the SP for case C2 let it be $[\hat{L}_b, \hat{L}_b]$. Also, let UP^a (UP^b) be a 's (b 's) approximate cumulative utility from $[L_a, L_b]$, \overline{UP}^a (\overline{UP}^b) be a 's (b 's) approximate cumulative utility from $[\bar{L}_a, \bar{L}_b]$, and \hat{UP}^a (\hat{UP}^b) be a 's (b 's) approximate cumulative utility from $[\hat{L}_b, \hat{L}_b]$. And let SW be the approximate social welfare for the PDP, \bar{SW} that for the case C1 for the SP, and \hat{SW} for case C2. In other words, we have $SW = UP^a + UP^b$, $\bar{SW} = \overline{UP}^a + \overline{UP}^b$, $\hat{SW} = \hat{UP}^a + \hat{UP}^b$.

For a given $m, n, \delta, \mathcal{A}_c$, and \mathcal{B}_c ($1 \leq c \leq m$), let P_{UA} (P_{UB}) denote the percentage difference in a 's (b 's) utilities for the PDP and case C1 for the SP. Likewise, let P_{SW} denote the percentage difference in the social welfare for the PDP and case C1 for the SP. Thus, for a given $m, n, \delta, \mathcal{A}_c$, and \mathcal{B}_c , we have:

$$P_{UA} = 100 \times (UP^a - \overline{UP}^a) / UP^a \quad (12)$$

$$P_{UB} = 100 \times (UP^b - \overline{UP}^b) / UP^b \quad (13)$$

$$P_{SW} = 100 \times (SW - \bar{SW}) / SW \quad (14)$$

For the case C2, \bar{P}_{UA} , \bar{P}_{UB} , and \bar{P}_{SW} are defined analogously.

The experimental setting is as follows. We vary the number of issues m between 50 and 100. For each m , we vary n (s.t. $n_\delta \leq n \leq 50$) in increments of 10. For each m and n , we vary δ between 0.1 and 0.9 in increments of 0.2. For each m, n , and δ , we randomly generate the weights (\mathcal{A}_c , and \mathcal{B}_c for $1 \leq c \leq m$) for the m issues using a uniform distribution in the interval $[1, 100]$. Here, the players are *symmetric* in that weights are independently and identically distributed across the two players. Note that the

probability distribution for the weights is identical for a and b but the actual weights are different for the two agents.

For each m, n , and δ , we find all six percentage differences (i.e., P_{UA} , P_{UB} , P_{SW} , \bar{P}_{UA} , \bar{P}_{UB} , and \bar{P}_{SW}) for 10 sets of randomly generated weights and find the average over these 10 cases. This is done for weights that are distributed uniformly on the interval $[1, 100]$. Let Δ_{UA} , Δ_{UB} , Δ_{SW} , $\bar{\Delta}_{UA}$, $\bar{\Delta}_{UB}$, and $\bar{\Delta}_{SW}$ denote the averages (taken over the 10 cases) for P_{UA} , P_{UB} , P_{SW} , \bar{P}_{UA} , \bar{P}_{UB} , and \bar{P}_{SW} respectively. It was found that each of the six percentage differences depended on m and δ but not on n^5 . Hence, we will show how the percentage differences vary with m and δ .

As mentioned before, the percentage differences we are measuring are with respect to the linear approximations of nonlinear utility functions. Thus, there is some error in these measurements which we quantify as follows. Recall from Section 4, that ϵ denotes an upper bound on the absolute error in an agent's cumulative utility computed using $\bar{F}A$ or $\bar{F}B$. Given ϵ , we find an upper bound on the error in the percentage differences (i.e., P_{UA} , P_{UB} , P_{SW} , \bar{P}_{UA} , \bar{P}_{UB} , or \bar{P}_{SW}) as follows. Consider P_{UA} (see Equation 12). In order to propagate the error (ϵ) to the error in P_{UA} , we use the following standard rules for error propagation [16]: If u and z are two random variables with errors ϵ_u and ϵ_z respectively, then the error in $u + z$ or $u - z$ is $\epsilon_u + \epsilon_z$. Also, the error in u/z is $\frac{\epsilon_u}{|u|} + \frac{\epsilon_z}{|z|}$, and the error in ku where k is a constant is $k\epsilon_u$.

Since the error in UP^a , UP^b , \overline{UP}^a or \overline{UP}^b is ϵ , the error in P_{UA} (denoted EP_{UA}) is $100 \times (\frac{2\epsilon}{|UP^a - \overline{UP}^a|} + \frac{\epsilon}{UP^a})$. In other words,

⁵This is because, as mentioned before, we vary n s.t. $n_\delta \leq n \leq 50$. From Figure 1, we know that, for $n_\delta \leq n \leq 50$, the players' shares for an issue do not vary much with n .

the true value of P_{UA} lies in the range $P_{UA} \pm EP_{UA}$. Likewise, the error in P_{UB} is $EP_{UB} = 100 \times (\frac{2\epsilon}{|UP^b - \bar{U}P^b|} + \frac{\epsilon}{UP^b})$. And the error in P_{SW} is $EP_{SW} = 100 \times (\frac{4\epsilon}{|SW - \bar{SW}|} + \frac{\epsilon}{SW})$. Finally, the error in \bar{P}_{UA} , \bar{P}_{UB} , or \bar{P}_{SW} is defined analogously.

Given these definitions, Tables 1 and 2 show a comparison of the PDP and the SP in a wide range of settings for cases C1 and C2. We first explain the results for the case C1 and then for C2. For the former, Table 1 shows Δ_{UA} , Δ_{UB} , and Δ_{SW} . For instance, for $m = 50$ and $\delta = 0.1$, $\Delta_{UA} = 7.98 \pm \epsilon$. Thus, which procedure is better for a depends on ϵ ; if $\epsilon < 7.98$, then Δ_{UA} is positive and therefore the PDP is better. On the other hand, for agent b , $\Delta_{UB} = 8.1 \pm \epsilon$, so the PDP is better if $\epsilon < 8.1$. In terms of social welfare, $\Delta_{SW} = 7.99 \pm 1.8\epsilon$ — so the PDP is better if $\epsilon < 4$. Likewise for other entries in the table.

Table 2 shows Δ_{UA} , Δ_{UB} , and Δ_{SW} for the case C2. In contrast to Table 1, some entries in this table are negative. For instance, for $m = 50$ and $\delta = 0.1$, $\Delta_{UB} = -503 \pm 0.3\epsilon$. Thus for $\epsilon < 1666$, the SP is better for b . On the other hand, for agent a , $\Delta_{UA} = 33.3 \pm 0.3\epsilon$, so the PDP is better if $\epsilon < 111$. Likewise for social welfare.

In summary, although the PDP is known to always generate Pareto optimal outcomes, the above analysis shows that the SP may be better for one of the two agents and also increase the social welfare.

6. RELATED WORK

Since Schelling first noted that the outcome of negotiation depends on the choice of negotiation procedure, much research effort has been devoted to the study of different procedures for negotiating multiple issues. This research can be broadly divided into two categories: negotiation with divisible issues [1, 4, 2] and with indivisible issues [14, 3]. For instance, [2] analyzes the equilibrium for the PDP and the SP in the context of divisible issues and linear utility functions. Again, in the context of divisible issues, [1, 4] examine the effect of negotiation agenda on the outcome of sequential negotiations. On the other hand, [14] focuses on indivisible issues and analyzes the task allocation problem when the agents maximize the benefit of the system as a whole. In contrast, our focus is on two agents where both of them are self-interested and want to maximize their individual utilities. In the same context, [3] shows that it is NP hard to find the equilibrium strategies for indivisible issues and linear utilities. However, a common feature of all the above work is that it deals with linear utility functions. In contrast, we focus on nonlinear utilities as a more realistic model in many settings.

A small number of researchers have tackled the problem of multi-issue negotiation with nonlinear utility functions under different simplifying assumptions. For instance, in [6], multi-issue negotiation is studied under the assumption that utility functions are nonlinear but are monotonic and concave. This assumption means that there is a unique global optimum which can easily be found (note that our work does not require monotonicity or concavity). On the other hand, in [7], the negotiating agents submit bids to a mediator which then chooses the final allocation. However, there is no game-theoretic analysis of the agents' strategic behavior and it is assumed that the agents will bid truthfully. Likewise, [8] deals with nonlinear utility functions but again there is no game-theoretic analysis that shows the proposed strategies form an equilibrium. In contrast, our work takes the strategic aspect into consideration and does a game-theoretic analysis of negotiation in terms of equilibrium outcomes. Also, in our work, the agents negotiate directly with each other without going through a mediator, which is, we believe, a more realistic setting for many problems.

In summary, the key differences between existing work and ours

are as follows. First, the former has dealt with linear utilities while we deal with nonlinear utilities. Second, much of the work has focused on a single negotiation procedure while we compare the equilibrium for two different procedures both in terms of the time complexities and approximation errors, and also in terms of the utilities they yield to the agents and the social welfare. Finally, negotiation for nonlinear utility functions has previously been studied, but under the assumption that agents do not behave strategically. In contrast, we take the strategic aspect into account and compare the two procedures in terms of their equilibrium outcomes.

7. CONCLUSIONS

This paper analyzes bilateral multi-issue negotiation between strategic self-interested agents. The issues are *divisible*, there are time constraints in the form of *deadlines* and *discount factors*, and the agents have different preferences over the issues. For such scenarios, we consider *nonlinear* utilities and show that finding the equilibrium for the PDP is computationally hard. Then, in order to overcome this complexity, we investigate two solutions: (i) approximating nonlinear utilities with linear ones and then using the PDP, and (ii) using the SP to negotiate the issues in parallel. This study shows that the equilibrium for the PDP (for approximate linear utilities) and the equilibrium for the SP (for nonlinear utilities) can be computed in time $\mathcal{O}(mn)$. In terms of the economic properties, although the PDP is known to generate Pareto optimal outcomes, we show that, in some cases, the SP may be better for one of the two agents and also increase the social welfare.

There are several interesting directions for future work. First, for the comparison of the two procedures, we focused on symmetric agents (i.e., for both agents, the weights were distributed identically). A comparison for asymmetric case will broaden the applicability of our results. Second, this work focused on separable nonlinear utility functions. An extension of the analysis to non-separable nonlinear utility functions is part of future work.

8. 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] S. S. Fatima and M. Wooldridge and N. R. Jennings, Multi-issue negotiation with deadlines, *Journal of Artificial Intelligence Research*, 27: 381–417, 2006.
- [3] S. S. Fatima and M. Wooldridge and N. R. Jennings, Approximate and online multi-issue negotiation, In *Proceedings of the Sixth International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-07)*, pp 947–954, 2007.
- [4] C. Fershtman, The importance of the agenda in bargaining, *Games and Economic Behavior*, 2: 224–238, 1990.
- [5] M. R. Garey and D. S. Johnson, *Computers and Intractability*, W. H. Freeman and Company, 1979.
- [6] F. Giunta and N. Gatti, Bargaining over multiple issues in finite horizon alternating offers protocol, *Annals of Mathematics and Artificial Intelligence*, 47: 251–271, 2006.
- [7] T. Ito and H. Hattori and M. Klein, Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces, In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pp 1347–1352, 2007.
- [8] M. Klein and P. Faratin and H. Sayama and Y. Bar-Yam, Negotiating complex contracts, *Group Decision and Negotiation*, 12(2): 58–73, 2003.
- [9] S. Kraus, Strategic negotiation in multi-agent environments, *The MIT Press, Cambridge, Massachusetts*, 2001.
- [10] S. Martello and P. Toth, *Knapsack problems: Algorithms and computer implementations*, John Wiley and Sons, 16–19, 1990.
- [11] R.E. Miller, *Optimization: Foundations and Applications*, Wiley-IEEE, 2000.
- [12] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, The MIT Press, 1994.
- [13] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter*, MIT Press, 1994.
- [14] O. Shehory and S. Kraus, Methods for task allocation via agent coalition formation, *Artificial Intelligence Journal*, 101(1-2), 165–200, 1998.
- [15] I. Stahl, *Bargaining Theory*, Economics Research Institute, Stockholm School of Economics, Stockholm, 1972.
- [16] J. R. Taylor, *An introduction to error analysis: The study of uncertainties in physical measurements*, University Science Books, 1982.