

Automated Multilateral Negotiation on Multiple Issues with Private Information

Ronghuo Zheng

McCombs School of Business, University of Texas at Austin, Austin, Texas 78705, ronghuo.zheng@mcombs.utexas.edu

Tinglong Dai

Carey Business School, Johns Hopkins University, Baltimore, Maryland, 21202, dai@jhu.edu

Katia Sycara[†]

Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, katia@cs.cmu.edu

Nilanjan Chakraborty

Department of Mechanical Engineering, Stony Brook University, Stony Brook, New York 11794, [nilanjan.chakraborty@stonybrook.edu](mailto:nilanjan.Chakraborty@stonybrook.edu)

In this paper, we propose and analyze a distributed negotiation strategy for a multi-agent, multi-attribute negotiation in which the agents have no information about the utility functions of other agents. We analytically prove that, if the zone of agreement is nonempty and the agents concede up to their reservation utilities, agents generating offers using our offer-generation strategy, namely the sequential projection strategy, will converge to an agreement acceptable to all the agents; the convergence property does not depend on the specific concession strategy. In considering agents' incentive to concede during the negotiation, we propose and analyze a reactive concession strategy. Through computational experiments, we demonstrate that our distributed negotiation strategy yields performance sufficiently close to the Nash bargaining solution and that our algorithms are robust to potential deviation strategies. Methodologically, our paper advances the state of the art of alternating projection algorithms, in that we establish the convergence for the case of multiple, moving sets (as opposed to two static sets in the current literature). Our paper introduces a new analytical foundation for a broad class of computational group decision and negotiation problems.

Keywords: operations research/computer science interfaces; convergence of automated negotiation; multi-agent, multi-attribute negotiation; alternating projection algorithms; distributed decision making

History: Accepted by Ram Ramesh, Area Editor for Knowledge Management and Machine Learning; received November 2014; revised October 2015; accepted February 2016. Published online August 17, 2016.

1. Introduction

Multi-attribute negotiation is a useful method in a wide range of scenarios, particularly when two or more parties (or agents) with limited common knowledge about each others' preferences try to arrive at an agreement on a set of issues over which they have possible conflicting preferences. For example, in international climate conferences, representatives of nations negotiate how much each nation will reduce carbon dioxide emissions (Lange et al. 2010). As another example, in e-commerce systems, firms are engaged in negotiating with multiple suppliers to fulfill orders (Zeng 2001). Our focus in this paper is multi-attribute automated negotiation among multiple agents from the perspective of designing systems of software agents to reach an agreement.

The extant literature on the mathematical study of negotiation can be divided into two broad categories:

mediated negotiations and nonmediated negotiations (Sycara and Dai 2010); in this paper, we focus on *nonmediated* negotiations. Whereas in a mediated negotiation, the presence of a nonbiased mediator is presumed and agents interact with each other through that mediator (Heiskanen et al. 2001, Klein et al. 2003, Ito et al. 2007, Chalamish and Kraus 2012), in a nonmediated negotiation, agents interact with each other directly.

In the nonmediated negotiation literature, researchers make different assumptions about the number of negotiating agents, the number of issues they are negotiating, and agents' knowledge regarding other agents' preferences (modeled using utility functions). Most work to date has focused on two-agent, single-issue negotiation, although some work has addressed two-agent, multi-issue negotiation (e.g., Fatima et al. 2006) or multi-agent, single-issue negotiation (e.g., Binmore 1985). Furthermore, computational modeling of multi-issue negotiation has either

[†]Corresponding author.

assumed (a) complete knowledge of the preference structure of the opponents, i.e., the utility functions of the agents are known, (e.g., Nash 1950, Rubinstein 1982), or (b) a probability distribution over the preferences of the agents is known (e.g., Harsanyi and Selten 1972, Chatterjee and Samuelson 1983, Lin et al. 2008). In addition, most of the literature assumes linear additive utility functions.

When the utility function is assumed to be linear and the information about the opponent’s utility function is known, a monotonic concession strategy and a Zeuthen strategy (Endriss 2006) have been proposed for negotiation. In the presence of incomplete information, Bayesian learning has been proposed in agents’ negotiation strategy (Li and Tesauro 2003, Buffett and Spencer 2005). Rational strategies that correspond to sequential equilibrium of a game have been proposed when each agent has probabilistic knowledge about its opponent (e.g., Fatima et al. 2004). However, these strategies cannot be used if knowledge about opponents’ utility functions is absent and when the utility functions are nonlinear. Preference elicitation—before or through negotiation—has been studied where agents have no knowledge about opponents’ utilities (e.g., Chari and Agrawal 2007). However, preference elicitation is known to be a difficult and time-consuming procedure (Chen and Pu 2004), especially when the agents’ preferences are complex. Most crucially, preference elicitation does not guarantee that an agreement will be reached even when the zone of agreement is nonempty.

In this paper, we study a multilateral negotiation in a general setting, where agents have nonlinear utility functions and aim to answer the following fundamental open question: *Is it possible to design a distributed negotiation strategy that enables agents to provably come to an agreement, given that they have no prior knowledge about the utility functions of other agents?* For agents with general utility functions, the utility of an offer is not simply a sum of the utilities of the individual issues. Therefore, we allow the agents to negotiate with package offers where agents negotiate on multiple issues simultaneously (as opposed to negotiating issue by issue). Although issue-by-issue offers are more convenient mathematically, packaged offers have the advantage that they allow agents to make trade-offs over different issues, which is a realistic feature of many negotiations. Our goal is to design a strategy for generating packaged offers for agents with private information that provably leads to an acceptable agreement for all the agents, providing the agents keep conceding to their reservation utilities and the zone of agreement is nonempty.

To illustrate the difficulty of generating an acceptable solution in a multilateral, multi-issue negotiation, let us consider the geometry of a negotiation

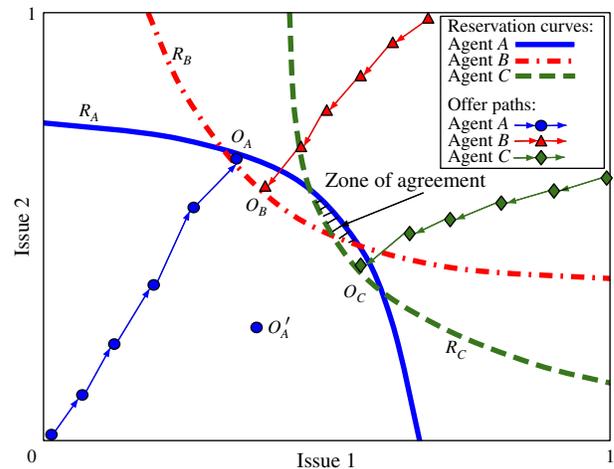


Figure 1 (Color online) Illustrative Sketch of the Offer Space of Three Agents, A, B, and C, Negotiating on Two Issues

Notes. R_A , R_B , and R_C are the reservation curves of the agents A, B, and C, respectively. The convex sets bounded by the three reservation curves are the feasible offer sets (e.g., O'_A is a feasible offer for Agent A).

problem. Figure 1 provides a geometric view of the offer space for three agents (referred to as A, B, C) negotiating on two issues. We denote by R_i ($i = A, B, C$) Agent i 's reservation curve, which is the set of all offers that provide the agent's reservation utility. In addition, we denote by O_i ($i = A, B, C$) an offer that Agent i chooses from R_i . For each agent, the convex set bounded by its reservation curve is the feasible offer set, and any offer within this set is acceptable to this agent because the utility of the offer to the agent is no less than the agent's reservation utility. For example, O'_A is a feasible offer for Agent A. The zone of agreement (the area in the common intersection or hatched region in Figure 1) is the set of offers that is acceptable to all agents. Any point within the zone of agreement is referred to as a *satisficing agreement*. Note that the zone of agreement is unknown to the agents because none of them knows any other agent's utility function. Thus, geometrically speaking, in negotiation, the goal of the agents is to find a point in the zone of agreement (i.e., a satisficing agreement), under the restriction that none of the agents has any explicit knowledge of the zone of agreement.

Let us step back and consider *two* agents negotiating on a *single* issue (e.g., a buyer and a seller negotiating on the price of a house). Here, if the zone of agreement is nonempty (i.e., if the lowest price at which the seller is willing to sell is less than the highest price the buyer is willing to pay), negotiation will always result in an agreement, because one agent's offer with utility equal to the agent's reservation utility would be acceptable to the other agent. Even for a multi-issue negotiation where the agents (with linear additive utility functions) negotiate issue by issue and

have a different reservation price¹ for each issue (e.g., Fatima et al. 2004 who consider two agents), an agreement can still be reached if every agent proposes an offer that corresponds to the agent's own reservation price for each issue.

For *packaged* multilateral, multi-attribute negotiation with nonlinear utility functions, however, it is nontrivial for an agent to find an offer acceptable to other agents. It is entirely likely that an offer on one agent's reservation utility curve is deemed unacceptable by another agent. Figure 1 shows that although the offers O_A , O_B , and O_C give the agents A , B , and C their lowest acceptable utilities (i.e., they concede as much as they can), these offers do not lie in the (unknown) zone of agreement, so none of them is acceptable to all the agents. In general, there are an infinite number of offers that lie on an agent's reservation utility curve that are unacceptable to the other agents. Thus, in contrast to single-issue negotiation, for multi-issue negotiation with private information, unknown zone of agreement, and nonlinear utility functions, it is challenging to generate offers acceptable to all agents in a distributed manner.

1.1. Contributions to the Literature

The first contribution we present is a distributed negotiation strategy for generating offers, referred to as *sequential projection strategy*; we analytically establish that the agents following this strategy will reach an agreement, assuming they concede to their reservation utilities and the zone of agreement is nonempty. Note that although the agents will concede, the amount by which they concede (or the rule by which they decide on the amount to concede) is not specified. Thus, there is some degree of freedom in the choice of concession rule (or concession strategy). We show that the convergence holds for general concave utility functions as long as all the agents concede to their reservation utilities, irrespective of the specific concession strategy the agents adopt. The sequential projection strategy is a generalization of the alternate projection heuristic that was proposed in the literature for two agents (Lai and Sycara 2009, Wu et al. 2009). We also prove that agents have no incentive to deviate from the sequential projection strategy for generating offers.

The second contribution is that this is the first attempt to tackle the issue of automated negotiating agents' incentives to concede. The literature has not explicitly discussed agents' incentives to concede in negotiations. One possible rationale in the literature for the lack of explicit accounting for incentives to concede is that the utility that agents obtain from reaching an agreement decreases with time. In other

words, the value of an outcome may be time sensitive and may decrease with time. However, in many negotiation scenarios, negotiators do concede with time, even in negotiations where the utility of issues does not decrease with time. This is because the agents desire to reach an agreement and know that if others see that they do not concede, then there is a chance that the negotiation may stall and one or more parties may walk out. Here we propose and analyze a concession strategy for negotiation that conforms to this intuition. In other words, the strategy is *reactive*; that is, it depends not only on whether other agents' offers give an agent utility higher than its reservation utility, but also on the agent's perception of how much others have conceded. Having the agents be reactive is novel. A common feature of concession strategies in the extant literature is that they are assumed to be exogenous and not reactive to the concession strategy of the other agents (notable exceptions being Aknine et al. 2004, Shakun 2005, Chari and Agrawal 2007). Shakun (2005) proposes a reactive tit-for-tat negotiation strategy and Aknine et al. (2004) propose an extension of contract net protocols to negotiations. In contrast, we design a reactive strategy and show that during a negotiation concession is rational even in the absence of time-decreasing utility functions. In particular, we prove that if the agents follow our reactive concession strategy, none of them will have an incentive to initiate stopping of concession.

To the best of our knowledge, for nonmediated negotiation, this paper is the first to provide a negotiation strategy with guaranteed convergence to a satisfying solution for general multi-attribute, multilateral negotiation with agents that have nonlinear utility functions and no knowledge about other agents' preferences. Moreover, we believe this paper is the first in negotiation literature that studies the issue of incentive compatibility of making concessions. We demonstrate through extensive simulations the performance of our algorithms and their robustness to noisy bids, among other deviation strategies.

A third contribution of our paper is that it advances the mathematics/computer science literature on alternating projection algorithms (Cheney and Goldstein 1959) in the following aspects: (a) we allow the sets to *move* over time and are the first to prove the convergence property under this novel setting, and (b) we allow *multiple* sets. The alternating projection algorithm that finds a point in the intersection of two closed convex sets by iteratively projecting a point first onto one set and then onto the other has been rediscovered many times in the literature (e.g., Bauschke and Borwein 1996, Combettes 1997). However, none of these papers allows the sets to move during the projecting as we do in this paper.

¹ Note that we use "price" here to be consistent with the literature (Fatima et al. 2004).

The remainder of this paper is organized as follows. In Section 2 we outline the framework of automated negotiation and formally state our research questions. In Section 3 we present our convergence proof of the proposed sequential projection strategy. The results in Section 3 do not depend on the specific implementation of the concession strategy. In Section 4 we provide a discussion of automated negotiating agents' incentive to concede; we prove it is rational for the agents to not deliberately stop conceding under our proposed reactive concession strategy. In Section 5 we present the results of our computational experiments on randomly generated negotiation instances and demonstrate our strategy yields performance sufficiently close to that of the Nash bargaining solution; we also demonstrate the robustness of our proposed strategies under potential deviating scenarios. Finally, in Section 6, we summarize our contributions and outline avenues for future research.

2. Negotiation Framework

We consider m self-interested agents $i \in \{1, 2, \dots, m\}$ negotiating on a set of issues $j \in \{1, 2, \dots, N\}$ with a time horizon of T periods. Let $[0, 1]$ denote the unit interval in \mathbb{R} and $[0, 1]^N$ be the unit hypercube in \mathbb{R}^N . Without loss of generality, we assume the issues take on continuous values and the negotiation domain for each issue is $\Omega_j = [0, 1]$ with 0 and 1 corresponding to the extreme values of the issues. Any point within the unit hypercube is a *package offer* or simply an *offer*. We assume Agent i 's utility function, $u_i(x)$, $i = 1, 2, \dots, m$, is continuous and concave $\forall x \in [0, 1]^N$. Without loss of generality, we can normalize the range of Agent i 's utility function to $[0, 1]$. The assumption that each agent's utility lies between 0 and 1 is not required for our main results to hold and is made purely for simplicity of presentation; the scale of the utility of each agent is of no critical importance, as long as the reservation utility and the scale of concession is consistent with the scale of the utility. We assume that for all agents, reaching an agreement has higher utility than negotiation breakdown with no agreement. We define the following concepts to formalize our negotiation framework.

DEFINITION 1. Agent i , $i \in \{1, 2, \dots, m\}$, has a *reservation utility*, ru_i , such that any offer with utility less than its reservation utility is not acceptable to Agent i .

DEFINITION 2. The *feasible offer set* of Agent i , $i \in \{1, 2, \dots, m\}$, denoted as A^i , is defined as the set of offers that provide Agent i with utility no less than Agent i 's reservation utility ru_i , namely, $A^i = \{x \in [0, 1]^N \mid u_i(x) \geq ru_i\}$.

DEFINITION 3. The *zone of agreement* \mathcal{Z} is the common intersection of the feasible offer sets of all agents, namely, $\mathcal{Z} = \bigcap_{i=1}^m A^i$.

Agent i 's feasible offer set A^i is strictly convex for each i . Because the zone of agreement is the intersection of a finite number of convex sets, it is also a convex set. For a negotiated agreement to exist, the zone of agreement must be nonempty. Any point within the zone of agreement is acceptable to every agent, and we call such a solution a *satisficing solution* to the negotiation.² Note that the zone of agreement is fixed given the utility functions and reservation utilities and does not change during the course of a negotiation.

2.1. Negotiation Protocol

A key issue in designing negotiation of software agents is to choose a negotiation protocol. For a two-agent negotiation, we will assume the agents use an *alternating-offer protocol* (Rubinstein 1982), where an agent proposes an offer and the other agent responds by either accepting it or proposing a new offer. For general multi-agent negotiations, we will use a generalization of the alternating-offer protocol, namely, a *sequential-offer protocol*, by which each agent proposes an offer in a predetermined sequence. We use a sequential protocol in the multi-agent setting and assume the agents propose their offers in a given order. The negotiation ends when an offer on the table is acceptable to all the agents (the notion of an *acceptable offer* during the negotiation will be defined precisely in the next section), or when the agents cannot find an offer acceptable to all (after some predetermined time period T). Next, we formally state our research problem.

Problem Statement: Given m agents negotiating on N issues where each agent has a strictly concave private utility function and the zone of agreement has a nonempty interior, find a concession strategy such that the agents have an incentive to follow the strategy and an offer-generation strategy such that it is guaranteed that the agents following the strategy will reach an agreement.

² Different definitions have been proposed for a *proper* negotiation solution. Axiomatic solution concepts have been proposed for bargaining games, which include the Nash bargaining solution (Nash 1950), the Kalai-Smorodinsky solution (Kalai and Smorodinsky 1975), the egalitarian solution (Kalai 1977), and the pareto-optimal solution. The set of points that satisfy these different solution requirements consists of subsets of the zone of agreement. However, computing the points requires that all the agents know each others' utility functions. Because an agent does not know the utility function of its opponent, we use a satisficing solution as our solution concept. A satisficing solution is any agreement that gives the negotiators a utility greater than or equal to their reservation utility. The use of a satisficing solution in this very general setting where the agents have no information about their opponents is in the spirit of Herbert Simon (1956). The control theory literature has also studied similar ideas and concepts (e.g., Stirling 2003, 2005; Lopes de Lima et al. 2015).

2.2. Overall Negotiation Strategy

For each agent, the proposed overall negotiation strategy consists of the consecutive application of the following two strategies: a concession strategy and an offer-generation strategy. When it is Agent i 's turn to make an offer, the agent first checks whether the current offer on the table (made by some other agent) is acceptable. If it is not, Agent i uses the concession strategy to determine its own current desirable concession utility and uses the offer-generation strategy to generate a new offer.

Concession Strategy. We first define an agent's indifferent surface (or curve).

DEFINITION 4. The set of offers that give Agent i a particular utility $u_i(t)$ is referred to as Agent i 's *indifference surface* (or *curve*) in period t .

During a negotiation, agents gradually reduce the utility of offers acceptable to them. A negotiating agent not only desires to reach an agreement with the other agents, but also wants to obtain as much utility as possible. Thus, when agents begin a negotiation, they propose offers generating the highest possible utility for themselves and gradually move toward offers generating lower utility. However, they will neither propose nor accept any offer with utility lower than their reservation utility.

DEFINITION 5. Agent i 's *desirable utility* in period t is $s_i(t)$, such that Agent i only accepts an offer that provides utility equal to or higher than $s_i(t)$ in period t .

DEFINITION 6. The set of offers that give Agent i desirable utility $s_i(t)$ is referred to as Agent i 's *concession surface* (or *curve*) in period t .

Note a concession surface (curve) is an indifference surface (curve) for Agent i , in a given time period t , but not vice versa.

DEFINITION 7. Agent i 's *concession strategy* is defined as a time series of the agent's desirable utility at time $0, 1, \dots$, namely, $(s_i(0), (s_i(1), s_i(2), \dots,))$; $s_i(t)$ is a *monotonically decreasing* function of t and $s_i(t) \geq ru_i, \forall t$.

Each agent uses the concession strategy at each period t to determine the desirable utility (and corresponding concession surface).

DEFINITION 8. For Agent i , let A_t^i be the set of all offers (including other agents' offers) that have utilities higher than $s_i(t)$ in period t . The set $A_t^i = \{x \in [0, 1]^N \mid u_i(x) \geq s_i(t)\}$, is referred to as Agent i 's *desirable offer set* in period t .

DEFINITION 9. An offer x_t at time t is referred to as an *acceptable offer* to Agent i if $x_t \in A_t^i$, that is, if the offer belongs to the agent's desirable offer set in period t .

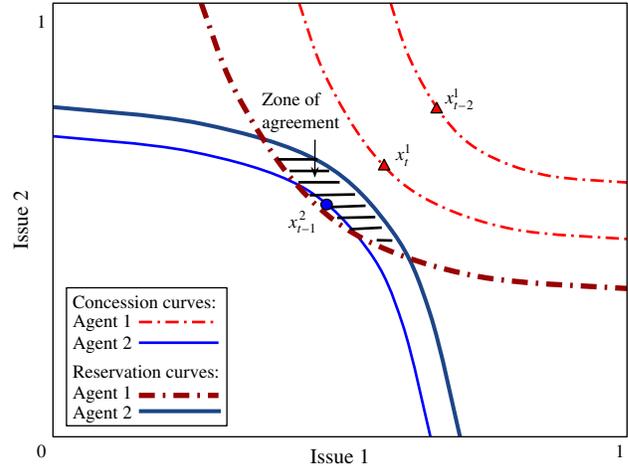


Figure 2 (Color online) Illustration of Definitions

For all t , as A_t^i represents the set of desirable offers in period t and Agent i keeps conceding over time, the desirable offer set keeps expanding over time; i.e., $A_t^i \subseteq A_{t+1}^i \subseteq \dots \subseteq A^i$. In Figure 2, e.g., the offer x_{t-2}^1 is on the desirable indifference curve for $t-2$ for Agent 1. All offers to the right and up from x_{t-2}^1 belong to Agent 1's current desirable offer set. At time t , as Agent 1 concedes, his current desirable offer set consists of all the points to the right and up from the current concession curve on which x_t^1 lies. Note Agent 1's desirable offer set at $t-2$ is a subset of its desirable offer set at time t .

From the prior discussion, an acceptable offer is always a feasible offer, but a feasible offer may not be an acceptable offer for the current period. Even if an offer is within the zone of agreement, an agent may not accept it when the offer yields a utility level that is below the agent's current desirable utility. As an example, in Figure 2, the offer x_{t-1}^2 is within the zone of agreement and a feasible offer for Agent 1 but not an acceptable offer for Agent 1 in period t .

Offer-Generation Strategy. The offer-generating strategy uses alternating projections. Here we define projection and the projection operator.

DEFINITION 10. For a convex set A and a point x , let $P_A[x]$ be the projection of point x on the set A with P being the projection operator. The projection $P_A[x]$ is a point in the set A that has the minimum Euclidean distance to x (Boyd and Vandenberghe 2004); e.g.,

$$P_A[x] = \arg \min_{q \in A} \|q - x\|, \quad (1)$$

where $\|\cdot\|$ denotes the Euclidean norm.

This definition implies that if $x \in A$, then $P_A[x] = x$. In other words, the projection of a point inside a set A on the same set A is the point itself. If $x \notin A$ and A is a compact set with its boundary defined by a differentiable function, then $P_A[x]$ lies on the boundary of the set A and the line joining x , and $P_A[x]$ is normal

to A .³ That is, $P_A[x]$ is the foot of the perpendicular from x on the boundary of A .

As previously stated, the agents make their proposals in a predetermined sequence. We assume at period $t=0$ that each agent proposes a utility-maximizing offer. After initialization, the agents propose sequentially, such that at time $t=1$, Agent 1 proposes, and, if the offer is not acceptable, then at time $t=2$, Agent 2 proposes and so on.

DEFINITION 11. An Agent j 's standing offer in period t , denoted by x_t^j , is the last offer Agent j made.

DEFINITION 12. Sequential Projection Strategy for Offer Generation: Let $t+1$ be the time when it is Agent $i+1$'s turn to propose. Agent $i+1$ determines the offer by projecting the convex combination of all the agents' standing offers (including Agent $i+1$'s previous offer) to Agent $t+1$'s current indifference surface in period $t+1$. More specifically, Agent $i+1$ at period $t+1$ proposes

$$x_{t+1}^{i+1} = P_{A_{t+1}^{i+1}} \left[\sum_{j=1}^m a_t^{i,j} x_t^j \right], \quad (2)$$

where A_{t+1}^{i+1} is Agent $i+1$'s set of acceptable offers at time $t+1$, $a_t^{i,j}$ is the weight that Agent i puts on Agent j 's standing offer at time t , and $\sum_{j=1}^m a_t^{i,j} = 1$.

To understand the working of the sequential projection strategy (referred to as alternating projection strategy for two agents), we start by examining the two-agent two-issue case (see Figure 3). In the figure, the dashed concession curves belong to Agent 1 and the solid concession curves belong to Agent 2. In period $t-4$, Agent 1 proposes an offer x_{t-4}^1 . Agent 2 rejects this offer and both agents update their indifference curves using their concession strategies. In period $t-3$, Agent 2 selects x_{t-3}^2 on the agent's concession curve such that x_{t-3}^2 is the projection of x_{t-4}^1 on Agent 2's concession curve. Agent 2 offers x_{t-3}^2 , which Agent 1 rejects, and both agents update their concession curves. In period $t-2$, Agent 1 identifies x_{t-2}^1 by projection of x_{t-3}^2 to Agent 1's current concession curve proposes it to Agent 2, and the negotiation proceeds.

Algorithm 1 (Overall algorithm)

Data: Each agent's utility function $u_i(x)$, reservation utility ru_i , and concession strategy $s_i(t)$, $t=1, 2, \dots, T$

Result: Negotiation Agreement

- 1 Initialization: Each agent proposes a preferred offer x_0^i ;

³ More generally, if the function defining the boundary is not differentiable everywhere, the line joining x and $P_A[x]$ lies in the normal cone to A at $P_A[x]$.

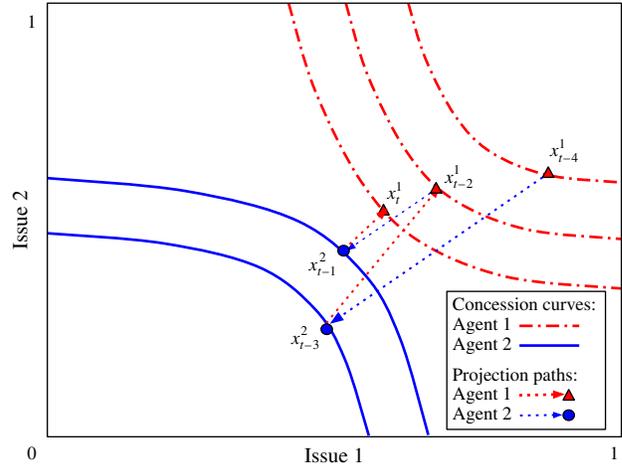


Figure 3 (Color online) Alternating Projection Strategy

```

2  $t = 1$ ;
3 Set convergence tolerance:  $\delta$ 
4 while  $t \leq T$  and  $IsConverge = False$  do
5     Determine the agent to propose:  $i = \text{mod}(t, m)$ ;
6     for each  $j \in \{1, 2, \dots, m\}$  do
7         if  $j = i$  then
8             Agent  $i$  concedes by determining  $s_i(t)$ ;
9             Agent  $i$  calculates:  $w_{t-1} \leftarrow (1/m) \sum_{j=1}^m x_{t-1}^j$ ;
10            Agent  $i$  proposes  $x_t^i(t) \leftarrow P_{A_i^i}[w_{t-1}]$ ;
11        else
12             $x_t^j \leftarrow x_{t-1}^j$ ;
13        end
14    end
15    if  $\max_{j \in \{1, 2, \dots, m\}} \|x_t^j - w_{t-1}\| < \delta$  then
16         $IsConverge = True$ 
17    else
18         $t = t + 1$ ;
19    end
20 end.
    
```

For general multilateral negotiation, to simplify the notation and presentation, we use $a_t^{i,j} = 1/m$ for $i, j \in \{1, 2, \dots, m\}$, $t \geq 0$; the discussion and results hold for general values of $a_t^{i,j}$ satisfying $\sum a_t^{i,j} = 1$. Therefore, Agent $i+1$ at period $t+1$ proposes an offer according to $x_{t+1}^{i+1} = P_{A_{t+1}^{i+1}}[(1/m) \sum_{j=1}^m x_t^j]$, where x_t^j is the newest offer proposed by Agent j until period t . For notational convenience, we define $w_t := (1/m) \sum_{j=1}^m x_t^j$. In other words, an agent selects an offer from the current concession surface by projecting the mean of all the other agents' latest offers to the agent's own current concession surface. Note this method generates an offer that is acceptable to the agent and is closest (in terms of Euclidean distance) to the average offer of the latest offers made by all agents.

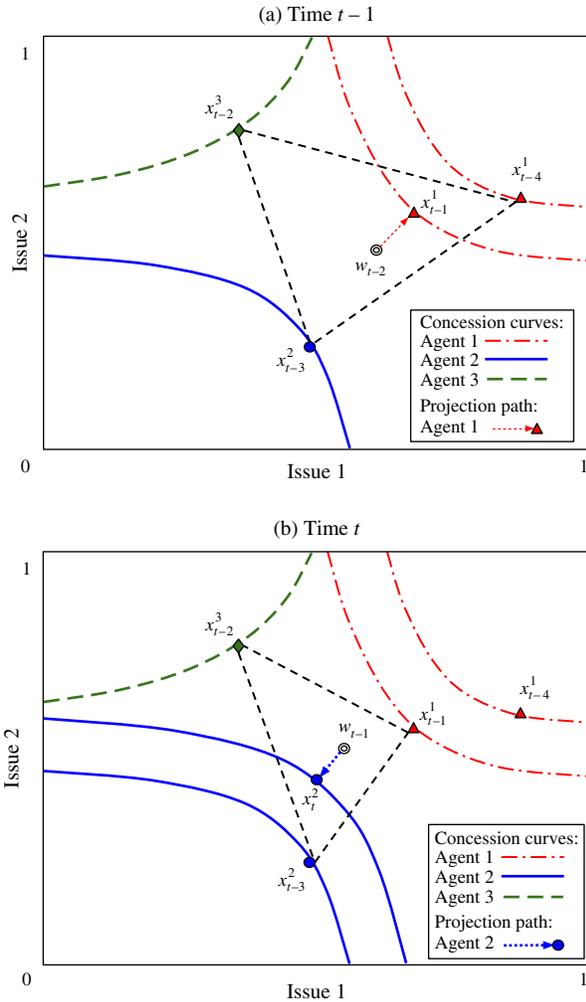


Figure 4 (Color online) The Sequential Projection Strategy for Three Agents Negotiating on Two Issues

Figure 4 illustrates the sequential projection method for three agents negotiating on two issues. At time $t - 1$ (see left-hand panel of Figure 4), it is Agent 1's turn to make an offer. The standing offers at period $t - 1$ are Agent 2's standing offer x_{t-3}^2 , because Agent 2's previous offer was proposed at time $t - 3$, and Agent 3's standing offer x_{t-2}^3 because Agent 3's previous offer was proposed at $t - 2$. Agent 1 incorporates the most recent offers from all agents, including the agent's own previous offer, namely, $x_{t-4}^1, x_{t-3}^2, x_{t-2}^3$, to compute the point w_{t-2} and project it onto Agent 1's concession surface at $t - 1$ (the dotted curve that is obtained using Agent 1's concession strategy). Similarly, in period t , it is Agent 2's turn to make an offer (see right-hand panel of Figure 4). Agent 2 computes the new offer by projecting the point w_{t-1} (computed by averaging $x_{t-1}^1, x_{t-3}^2, x_{t-2}^3$) to get x_t^2 , and the negotiation proceeds. In Algorithm 1, we provide the pseudocode for the overall negotiation strategy.

3. Convergence of the Sequential Projection Strategy

In this section, we establish the convergence of the sequential projection strategy for multilateral, multi-issue negotiation. We prove the convergence result for any concession strategy, without restricting ourselves to the specific implementation of the concession strategy. This implies that our results hold even if each agent in the negotiation has a different concession strategy. Under such a general setting, we show our sequential projection negotiation strategy ensures the agents converge to an agreement. We then examine the question of whether agents following the sequential projection strategy will find an agreement in a finite time.

3.1. Reaching Agreement with the Sequential Projection Strategy

We first restate a classical result from the convex geometry literature (Cheney and Goldstein 1959). Let $P_A[x]$ be the projection of point x on the set A with P being the projection operator, and let the notation $\|\cdot\|$ denote the Euclidean norm of a vector.

LEMMA 1 (CHENEY AND GOLDSTEIN 1959). Let A be a nonempty closed convex set in $[0, 1]^N$. Then $\forall x \in [0, 1]^N, y \in A$, and we have the following:

$$(P_A[x] - y)'(y - x) \leq -\|P_A[x] - y\|^2, \quad (3)$$

$$\|P_A[x] - x\|^2 \leq \|x - y\|^2 - \|P_A[x] - y\|^2. \quad (4)$$

Lemma 1 supports the convergence for alternating projection between *two static* sets. In Theorem 1, by contrast, we establish the convergence for alternating projection between *multiple moving* sets.

THEOREM 1. Let x_t^i be the latest offer proposed by Agent i until period t , and let w_t be the mean of the standing offers from all agents in period t . Then the sequence $\{\sum_{i=1}^m \|x_t^i - w_t\|^2\}$ is nonincreasing with t .

PROOF. Let Agent $i + 1$ be the agent proposing an offer x_{t+1}^{i+1} in period $t + 1$. Then we have

$$x_{t+1}^{i+1} = P_{A_{t+1}^{i+1}}[w_t]. \quad (5)$$

Because $x_{t+1}^j = x_t^j, j \neq i + 1$, we have

$$\begin{aligned} w_{t+1} &= \frac{1}{m} \sum_{j=1}^m x_{t+1}^j \\ &= \frac{1}{m} \left(\sum_{j=1}^m x_t^j + x_{t+1}^{i+1} - x_t^{i+1} \right) \\ &= w_t + \frac{1}{m} (x_{t+1}^{i+1} - x_t^{i+1}). \end{aligned} \quad (6)$$

Using the prior results, we have

$$\begin{aligned}
 & \sum_{j=1}^m \|x_{t+1}^j - w_{t+1}\|^2 \\
 &= \sum_{j=1, j \neq i+1}^m \|x_{t+1}^j - w_{t+1}\|^2 + \|x_{t+1}^{i+1} - w_{t+1}\|^2 \\
 &= \sum_{j=1, j \neq i+1}^m \left\| (x_t^j - w_t) - \frac{1}{m} (x_{t+1}^{i+1} - x_t^{i+1}) \right\|^2 \\
 &\quad + \left\| (x_{t+1}^{i+1} - x_t^{i+1}) + (x_t^{i+1} - w_t) - \frac{1}{m} (x_{t+1}^{i+1} - x_t^{i+1}) \right\|^2 \\
 &= \sum_{j=1}^m \left\| (x_t^j - w_t) - \frac{1}{m} (x_{t+1}^{i+1} - x_t^{i+1}) \right\|^2 \\
 &\quad + \|x_{t+1}^{i+1} - x_t^{i+1}\|^2 - \frac{2}{m} \|x_{t+1}^{i+1} - x_t^{i+1}\|^2 \\
 &\quad + 2(x_{t+1}^{i+1} - x_t^{i+1})'(x_t^{i+1} - w_t), \tag{7}
 \end{aligned}$$

in which

$$\begin{aligned}
 & \sum_{j=1}^m \left\| (x_t^j - w_t) - \frac{1}{m} (x_{t+1}^{i+1} - x_t^{i+1}) \right\|^2 \\
 &= \sum_{j=1}^m \|x_t^j - w_t\|^2 + \frac{1}{m} \|x_{t+1}^{i+1} - x_t^{i+1}\|^2 \\
 &\quad + \frac{2}{m} (x_{t+1}^{i+1} - x_t^{i+1})' \sum_{j=1}^m (x_t^j - w_t). \tag{8}
 \end{aligned}$$

By the definition of w_t , we have $\sum_{j=1}^m (x_t^j - w_t) = 0$. Therefore, we have

$$\begin{aligned}
 \sum_{j=1}^m \|x_{t+1}^j - w_{t+1}\|^2 &= \sum_{j=1}^m \|x_t^j - w_t\|^2 + \frac{m-1}{m} \|x_{t+1}^{i+1} - x_t^{i+1}\|^2 \\
 &\quad + 2(x_{t+1}^{i+1} - x_t^{i+1})'(x_t^{i+1} - w_t). \tag{9}
 \end{aligned}$$

We have from Lemma 1 that

$$(x_{t+1}^{i+1} - x_t^{i+1})'(x_t^{i+1} - w_t) \leq -\|x_{t+1}^{i+1} - x_t^{i+1}\|^2, \tag{10}$$

which in turn gives

$$\begin{aligned}
 \sum_{j=1}^m \|x_{t+1}^j - w_{t+1}\|^2 &\leq \sum_{j=1}^m \|x_t^j - w_t\|^2 - \frac{m+1}{m} \|x_{t+1}^{i+1} - x_t^{i+1}\|^2 \\
 &\leq \sum_{j=1}^m \|x_t^j - w_t\|^2. \quad \square
 \end{aligned}$$

The implications of Theorem 1 are twofold. First, in terms of convex geometry, Cheney and Goldstein (1959) deal with static sets, whereas Theorem 1 deals with moving sets. Second, in terms of negotiation, Theorem 1 states that the sequential projection strategy ensures the distance between the new offer generated by an agent and the mean of all the previous offers never increases. We use this fact to prove the next theorem, which is one of our main contributions.

THEOREM 2. *If the zone of agreement has a nonempty interior, and if the agents keep conceding to their reservation utilities, then the sequential projection proposing strategy will always converge to an agreement.*

PROOF. Suppose there exists some s such that $A_s = \bigcap_{i=1}^m A_s^i \neq \emptyset$ (otherwise, $\forall t, \bigcap_{i=1}^m A_t^i = \emptyset$, which implies the set $\lim_{t \rightarrow \infty} \bigcap_{i=1}^m A_t^i$ contains no interior point, which contradicts our assumption that the zone of agreement is nonempty). Then, as $A_t^i \subset A_{t+1}^i$ for $i \in \{1, 2, \dots, m\}$, we have $\forall t \geq s, \bigcap_{i=1}^m A_t^i \neq \emptyset$. Let

$$e_t^i = P_{A_t^i}[w_{t-1}] - w_{t-1}. \tag{11}$$

Without loss of generality, we assume Agent 1 proposes an offer x_{s+1}^1 in period $s+1$. Because we follow the convention that Agent 1 proposes in period $s+1$, Agent 2 proposes in period $s+2$, and so on, we have that Agent i proposes an offer x_{s+i}^i in period $s+i$. We also get the implication that Agent i proposes offers in all subsequent periods of the form $s+k \cdot m+i$, $\forall i \in \{1, 2, \dots, m\}$ and $\forall k \in \mathbb{N}$. Then, by Lemma 1, $\forall i \in \{1, 2, \dots, m\}$, and $\forall x \in A_s$, we get

$$\begin{aligned}
 \|e_{s+i}^i\|^2 &\leq \|w_{s+i-1} - x\|^2 - \|x_{s+i}^i - x\|^2 \\
 &= \left\| \frac{1}{m} \sum_{j=1}^m (x_{s+i-1}^j - x) \right\|^2 - \|x_{s+i}^i - x\|^2 \\
 &\leq \left(\frac{1}{m} \sum_{j=1}^m \|x_{s+i-1}^j - x\| \right)^2 - \|x_{s+i}^i - x\|^2 \\
 &\leq \frac{1}{m} \sum_{j=1}^m \|x_{s+i-1}^j - x\|^2 - \|x_{s+i}^i - x\|^2, \tag{12}
 \end{aligned}$$

where x_{s+i}^i is the latest offer proposed by Agent i until period $s+i$.

Thus, by summing (12) over all agents, we get

$$\sum_{i=1}^m \|e_{s+i}^i\|^2 \leq \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^m \|x_{s+i-1}^j - x\|^2 - \sum_{i=1}^m \|x_{s+i}^i - x\|^2. \tag{13}$$

Moreover, note that x_s^i is the latest offer proposed by Agent i until period s . Thus, we have $x_s^i = x_{s+1}^i = \dots = x_{s+i-1}^i$ and $x_{s+i}^i = x_{s+i+1}^i = \dots = x_{s+m}^i$. Therefore,

$$\begin{aligned}
 & \sum_{i=1}^m \sum_{j=1}^m \|x_{s+i-1}^j - x\|^2 \\
 &= \sum_{i=1}^m \left(\sum_{j=1}^{i-1} \|x_{s+j}^j - x\|^2 + \sum_{j=i}^m \|x_{s+j-1}^j - x\|^2 \right) \\
 &= \sum_{j=1}^m \sum_{i=j+1}^m \|x_{s+j}^j - x\|^2 + \sum_{j=1}^m \sum_{i=1}^j \|x_{s+j-1}^j - x\|^2 \\
 &= \sum_{j=1}^m (m-j) \|x_{s+j}^j - x\|^2 + \sum_{j=1}^m j \|x_{s+j-1}^j - x\|^2. \tag{14}
 \end{aligned}$$

By substituting (14) into (13), we obtain

$$\sum_{i=1}^m \|e_{s+i}^i\|^2 \leq \sum_{i=1}^m \frac{i}{m} (\|x_{s+i-1}^i - x\|^2 - \|x_{s+i}^i - x\|^2). \quad (15)$$

The inequality (15) holds by replacing s with $s + km$, $\forall k \in \mathbb{N}$. Moreover, $\forall k \in \mathbb{N}$, $x_{s+km+i}^i = x_{s+(k+1)m+i-1}^i$. Therefore, $\forall r \in \mathbb{N}$,

$$\begin{aligned} & \sum_{k=0}^r \sum_{i=1}^m \|e_{s+km+i}^i\|^2 \\ & \leq \sum_{k=0}^r \sum_{i=1}^m \frac{i}{m} (\|x_{s+km+i-1}^i - x\|^2 - \|x_{s+km+i}^i - x\|^2) \\ & = \sum_{i=1}^m \frac{i}{m} \sum_{k=0}^r (\|x_{s+km+i-1}^i - x\|^2 - \|x_{s+(k+1)m+i-1}^i - x\|^2) \\ & = \sum_{i=1}^m \frac{i}{m} (\|x_{s+i-1}^i - x\|^2 - \|x_{s+(r+1)m+i-1}^i - x\|^2) \\ & \leq \sum_{i=1}^m \frac{i}{m} \|x_{s+i-1}^i - x\|^2. \end{aligned} \quad (16)$$

When $r \rightarrow \infty$, the inequality (16) implies

$$\lim_{k \rightarrow \infty} \sum_{i=1}^m \|e_{s+km+i}^i\|^2 = 0. \quad (17)$$

Hence, $\lim_{t \rightarrow \infty} \|e_t^i\| = 0$ for all i . \square

In the proof of Theorem 2, we first show that if the zone of agreement has a nonempty interior, the sum of the sequences $\{\sum_{i=1}^m \|x_t^i - w_t\|^2\}$ has a finite upper bound. Moreover, we have proven the nonincreasing nature of the sequences $\{\sum_{i=1}^m \|x_t^i - w_t\|^2\}$ in Theorem 1, so the sequence $\{\sum_{i=1}^m \|x_t^i - w_t\|^2\}$ converges to 0, which implies the sequential projection proposing strategy will always converge to an agreement.

3.2. Reaching Agreement in a Finite Time

We now address the following question: Given that the concession strategies of the agents are such that all the agents reach their reservation utilities within a finite time, say T_0 , can the agents converge to an agreement in a finite time (provided the zone of agreement is nonempty)? Note that, as in Section 3.1, we do not make any assumptions about the specific concession strategy used by the agents.

THEOREM 3. *For m agents negotiating on N issues, if the agents use concession strategies such that they each reach their reservation utilities in a finite time, they will reach an agreement in a finite time (assuming the zone of agreement has a nonempty interior).*

PROOF. From inequality (16) in Theorem 2, we have $\forall r \in \mathbb{N}$, where \mathbb{N} is the set of integers,

$$\begin{aligned} & \sum_{k=1}^r \sum_{i=1}^m \|e_{s+km+i}^i\|^2 \\ & \leq \sum_{i=1}^m \frac{i}{m} (\|x_{s+i-1}^i - x\|^2 - \|x_{s+r.m+i}^i - x\|^2). \end{aligned} \quad (18)$$

Moreover, we have from the definition of e_t^i that

$$\sum_{i=1}^m \|e_{s+(k+1)m+i}^i\|^2 \leq \sum_{i=1}^m \|e_{s+km+i}^i\|^2, \quad \forall k \in \mathbb{N}. \quad (19)$$

Thus,

$$\sum_{i=1}^m \|e_{s+r.m+i}^i\|^2 \quad (20)$$

$$\leq \frac{1}{r} \sum_{i=1}^m \frac{i}{m} (\|x_{s+i-1}^i - x\|^2 - \|x_{s+r.m+i}^i - x\|^2)$$

$$\leq \frac{1}{r} \sum_{i=1}^m \frac{i}{m} \|x_{s+i-1}^i - x\|^2, \quad (21)$$

which implies $\forall \varepsilon > 0$, there exists $\tau > 0$, where

$$\tau = s + r \left\lceil \frac{\sum_{i=1}^m (i/m) \|x_{s+i-1}^i - x\|^2}{\varepsilon} \right\rceil + i \quad (22)$$

such that $\forall t > \tau$, $\sum_{i=1}^m \|e_t^i\|^2 < \varepsilon$. \square

The intuition of the proof of Theorem 3 is as follows: because we can guarantee the convergence of the sequence $\{\sum_{i=1}^m \|x_t^i - w_t\|^2\}$ using the fact that the sequence is nonincreasing (by Theorem 1) and the fact that the sequence has a finite sum (shown in Theorem 2), we have $\forall \varepsilon > 0$, $\exists \tau > 0$, s.t., $\forall t > \tau$, $\sum_{i=1}^m \|x_t^i - w_t\|^2 < \varepsilon$, where x_t^i is Agent i 's offer in period t and w_t is the mean of the standing offers of all the agents at period t . In other words, the distance between the offers generated by an agent and the mean of the current offers of all the agents will decrease to zero (within a numerical error tolerance ε) in a finite time.

4. Incentive for Agents to Concede

As discussed in Section 1.1, the negotiation literature does not explicitly examine whether the agents have incentive to concede. We believe that our work is the first to tackle the subject of incentives for agents to concede. The framework for addressing this is based on two assumptions. First, the agent's utility from reaching agreement is higher than nonagreement. Therefore, the agents would prefer conceding over risking a negotiation breakdown. Second, the agents are reactive (as opposed to the literature, in which their concession rate is exogenously determined), so their concession rate depends on their perception of

the utility of other agents' offers. We design a *reactive concession strategy*, where an agent concedes according to two criteria: (a) whether the current offer of other agents provides higher utility than the agent's reservation utility and (b) the agent's perception of how much the other agents have conceded. The reason an agent may want to stop conceding—before reaching its reservation utility—is to gain higher utility, if the other agents accept its (nonconcession) offer without realizing the agent has stopped conceding. In multilateral, multi-issue negotiation, an agent may stop conceding before reaching its reservation utility and simply make “concessions” on its current indifference curve. Since such pseudoconcessions are difficult for other agents to perceive, the agent is in effect manipulating the others and unfairly gaining more utility. We call this behavior *deliberate stopping* of concession.

From the point of view of designing an incentive-compatible reactive strategy, the difficulty of other agents of perceiving a deliberate stopping of concession is one challenge. We would like to design our strategy so that agents could perceive the nonconcession and then stop conceding as well, so as not to be taken advantage of. However, we would also like to distinguish the case of deliberate stopping of concession from the case of a nonmanipulative agent who happens to arrive at its reservation utility very early in the negotiation and thus has no other choice but to concede only on the reservation curve. In this case, if agents perceive the nonconcession, then they may stop conceding and the negotiation may break down, although, assuming the zone of agreement was nonempty, an acceptable agreement would have been found; had all other agents continued conceding and thus reaching their reservation utilities, the sequential projection strategy for offer generation would guarantee convergence to agreement. This is a second challenge in designing an incentive-compatible reactive strategy. Therefore, besides the deliberative concession stopping, where an agent initiates the stopping, we define a *reactive concession stopping*, where an agent stops conceding in reaction to another agent's stopping to concede.

We design a *reactive concession strategy* and show that under such a strategy, agents have an incentive to concede. To have a well-defined conceptual framework that is in line with the spirit of incentive compatibility (Myerson 1979), we define the notion of weak incentive compatibility:

DEFINITION 13. A strategy satisfies *weak incentive compatibility* or is *weakly incentive compatible* if and only if no other strategy dominates the strategy.

4.1. Reactive Concession Strategy

According to Definition 7, Agent i 's concession strategy is defined as a time series of the agent's desirable utility. Under a nonreactive concession strategy, Agent i would concede following a predefined concession strategy $(s_i^0(1), s_i^0(2), \dots, s_i^0(T))$. A reactive concession strategy is represented by $(s_i(1), s_i(2), \dots, s_i(T))$. To determine the amount of utility to concede at period t , each agent considers its own utility change resulting from other agents' offers. There are two cases. First, the change in utility that other agents' offers caused has resulted in *higher* utility than Agent i 's reservation utility. In this case, Agent i 's utility is reduced according to the nonreactive concession strategy, i.e., $s_i(t) = s_i^0(t)$ or, equivalently, by conceding $\Delta u_{i0}(t) = s_i^0(t) - s_i(t - 1)$.

Second, suppose that in period t , the change in utility from other agents' offers has resulted in *lower* utility than Agent i 's reservation utility. Then Agent i concedes by an amount based on the change Agent i perceives in its own utility resulting from other agents' offers.

DEFINITION 14. Let $x_{[i,-1]}^j(t)$ be Agent j 's *next-to-last best offer*, which is the offer that provides the highest utility to Agent i among all offers made by Agent j until Agent j 's next-to-last offer (i.e., not including Agent j 's standing offer) in period t .

The marginal perceived change of utility for Agent i from Agent j 's standing offer $x_i^j(t)$, is defined as $\Delta_1 u_{ij}(t) = u_i(x_i^j(t)) - u_i(x_{[i,-1]}^j(t))$. The total perceived change of utility for Agent i from Agent j 's standing offer $x_i^j(t)$, is $u_i(x_i^j(t)) - u_i(x_0^j)$. The total concession by Agent i until period $t - 1$ is $1 - u_i(x_{t-1}^i)$. Denote $\Delta_2 u_{ij}(t) = u_i(x_i^j(t)) - u_i(x_0^j) - (1 - u_i(x_{t-1}^i))$ as the difference between the total perceived concession from Agent j 's standing offer for Agent i and the total concession by Agent i . We define the *reactive concession* of Agent i in response to Agent j in period t as $\Delta u_{ij}(t) = \max\{\Delta_1 u_{ij}(t), \Delta_2 u_{ij}(t), 0\}$. Denote $\Gamma_t(i) = \{j \mid u_i(x_t^j) \leq ru_i\}$ as the set of agents whose standing offers provide lower utility to Agent i than Agent i 's reservation utility in period t . Then the desirable utility Agent i is willing to concede in period t is given by

$$\Delta u_i(t) = \min \left\{ \min_{j \in \Gamma_t(i)} \Delta u_{ij}(t), \Delta u_{i0}(t) \right\}. \quad (23)$$

Thus, according to the reactive concession strategy, the desirable utility of Agent i in round t is given by $s_i(t) = s_i(t - 1) - \Delta u_i(t)$. We provide the pseudocode for the reactive concession strategy in Algorithm 2. In addition, we provide the pseudocode for the overall algorithm with reactive concession strategy in Algorithm 3.

Algorithm 2 (Reactive concession algorithm)

Data: Each agent's standing offer x_t^i , second latest offer $x_{[i,-1]}^j$; Agent i 's utility function $u_i(x)$, reservation utility ru_i , and non-reactive concession strategy $s_i^0(t)$, $t = 1, 2, \dots, T$

Result: Agent i 's desirable utility in period t $s_i(t)$.

```

1 Nonreactive concession:  $\Delta u_{i0}(t) \leftarrow s_i^0(t) - s_i(t-1)$ 
2 for each  $j \in \{1, 2, \dots, m\}$  do
3   if  $x_{[i,-1]}^j = \emptyset$  or  $u_i(x_t^i) \geq ru_i$  then
4      $\Delta u_{ij}(t) \leftarrow \Delta u_{i0}(t)$ 
5   else
6      $\Delta_1 u_{ij}(t) \leftarrow u_i(x_t^i) - u_i(x_{[i,-1]}^j)$ 
7      $\Delta_2 u_{ij}(t) \leftarrow u_i(x_t^i) - u_i(x_0^i) - (1 - u_i(x_{[i,-1]}^j))$ 
8      $\Delta u_{ij}(t) \leftarrow \max\{\Delta_1 u_{ij}(t), \Delta_2 u_{ij}(t), 0\}$ 
9   end
10  Determine the reactive concession:
11    $\Delta u_i(t) \leftarrow \min\{\min_{j \in \Gamma_i(t)} \Delta u_{ij}(t), \Delta u_{i0}(t)\}$ 
12  Revise the desirable utility in period  $t$ :
13    $s_i(t) \leftarrow s_i(t-1) - \Delta u_i(t)$ 
14 end.

```

4.2. Weak Incentive Compatibility to Concede

We first show that, *assuming every agent concedes, the sequential projection strategy is weakly incentive compatible*. This is true, because all points on Agent i 's concession surface give the agent the same utility, and by proposing any other point, the agent may decrease the chance of reaching an agreement (because the convergence proof holds only for projections).

We then tackle the issue of agents' incentive to concede. When Agent i uses the reactive concession strategy, in period t , Agent i may stop conceding (i.e., $\Delta u_i(t) = 0$), because either $\Delta u_{i0}(t) = 0$ or $\min_{j \in \Gamma_i(t)} \Delta u_{ij}(t) = 0$. We define the case where $\Delta u_{i0}(t) = 0$ as *deliberately ceasing to concede*⁴ and the case where $\min_{j \in \Gamma_i(t)} \Delta u_{ij}(t) = 0$ as *reactively ceasing to concede*. We show that if all other agents use the reactive concession strategy, not deliberately ceasing to concede is weakly incentive compatible for an agent.

As the first step of the proof, in Lemma 2, we show the negotiation will stall if an agent deliberately stops conceding under certain conditions. To illustrate the conditions, as shown in Figure 5, let Agent i propose x_t^i at time t . Let $x_{s_i(t)}^*$ be the point on the indifference surface $u_i(x) = s_i(t)$ such that $u_j(x) = u_j(x_{s_i(t)}^*)$ is the highest attainable utility by Agent j on this indifference surface. Let $x_{ru_i}^*$ be the point on Agent i 's reservation surface such that $u_j(x) = u_j(x_{ru_i}^*)$ is the highest attainable utility by Agent j on this indifference surface. Define $\Delta_j \equiv u_j(x_{s_i(t)}^*) - u_j(x_t^i)$.

⁴Note that by definition, $\Delta u_{i0}(t) = 0$ is Agent i 's concession when other agents' offers have given Agent i utility higher than the agent's own reservation utility; i.e., deliberate stopping of concession happens before Agent i reaches its own reservation utility.

Algorithm 3 (Overall algorithm with reactive concession strategy)

Data: Each agent's utility function $u_i(x)$, reservation utility ru_i , and planned concession strategy $s_i^0(t)$, $t = 1, 2, \dots, T$

Result: Negotiation Agreement

```

1 Initialization: Each agent proposes a preferred offer  $x_0^i$ , and sets  $x_{[i,-1]}^j \leftarrow \emptyset$ ,  $\forall j \in \{1, 2, \dots, i-1, i+1, \dots, m\}$ .
2  $t = 1$ 
3 Set convergence tolerance:  $\delta$ 
4 while  $t \leq T$  and  $IsConverge = False$  do
5   Determine the agent to propose:  $i = \text{mod}(t, m)$ ;
6   for each  $j \in \{1, 2, \dots, m\}$  do
7     if  $j = i$  then
8       Agent  $i$  concedes using Algorithm 2:
9        $s_i(t) \leftarrow s_i(t-1) - \Delta u_i(t)$ ;
10      Agent  $i$  calculates:  $w_{t-1} \leftarrow (1/m) \sum_{j=1}^m x_{t-1}^j$ ;
11      Agent  $i$  proposes  $x_t^i \leftarrow P_{A_i}[w_{t-1}]$ ;
12    else
13       $x_t^j \leftarrow x_{t-1}^j$ ;
14      if  $u_j(x_{t-1}^j(t-1)) \geq u_j(x_{[j,-1]}^i(t-1))$  then
15         $x_{[j,-1]}^i(t) \leftarrow x_{t-1}^i(t-1)$ ;
16      end
17       $x_{[i,-1]}^j(t) \leftarrow x_{[i,-1]}^j(t-1)$ ;
18    end
19  if  $\max_{j \in \{1, 2, \dots, m\}} \|x_t^j - w_{t-1}\| < \delta$  then
20     $IsConverge = True$ ;
21  else
22     $t = t + 1$ ;
23  end
24 end.

```

LEMMA 2. *If Agent i deliberately stops conceding before reaching the agent's own reservation utility from time period t onward, and all other agents use the reactive concession strategy, the negotiation will stall; i.e., other agents*

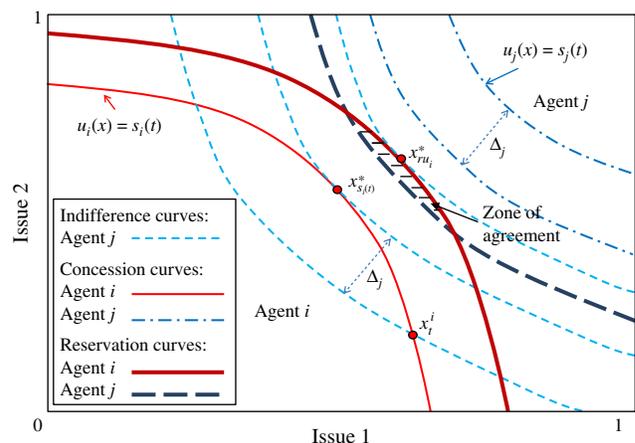


Figure 5 (Color online) Agents' Incentive to Concede

will reactively stop conceding, and there will be no agreement, if $\Delta_j < s_j(t) - u_j(x_{ru_j}^*)$ and $u_j(x_{s_i(t)}^*) < ru_j$.

PROOF. If Agent i stops conceding from period t onward, all offers that Agent i proposes after period t are on Agent i 's indifference surface $u_i(x) = u_i(x_i^t) = s_i(t)$. Therefore, from period $t + 1$ onward, the maximum of the total perceived utility improvement by Agent j from Agent i 's offers is

$$\Delta_j \equiv u_j(x_{s_i(t)}^*) - u_j(x_i^t). \quad (24)$$

Without loss of generality, from period t , the minimum of the total utility improvement Agent j perceives from the offers made by the other agents (including Agent i) would be smaller than or equal to the total utility improvement that Agent j perceives from Agent i 's offers. If Agent i stops conceding and $u_j(x_{s_i(t)}^*)$ remains smaller than Agent j 's reservation utility ru_j , using the reactive strategy, Agent j would concede by at most a total of Δ_j over the next rounds, irrespective of Agent i 's offers. Note that Agent j will not know the amount Δ_j , but the nature of the reactive concession strategy guarantees that Agent j 's total concession from t onward is no more than Δ_j . Thus, if $\Delta_j < s_j(t) - u_j(x_{ru_j}^*)$ and $u_j(x_{s_i(t)}^*) < ru_j$, where $s_j(t)$ is Agent j 's current utility level in period t , the negotiation will stall; that is, Agent j will stop conceding outside of the zone of agreement, and thus Agents i and j will not reach an agreement, although there is a nonempty zone of agreement. \square

THEOREM 4. If all the agents use the reactive concession strategy, none of them has an incentive to deliberately stop conceding; thus, it is weakly incentive compatible for all agents to keep conceding throughout the negotiation process.

PROOF. By Lemma 2, and because Agent i has no knowledge of other agents' utility functions, the agent is uncertain about whether $\forall j, j \in \{1, 2, \dots, m\}, j \neq i$, $u_j(x_{s_i(t)}^*)$ is higher than Agent j 's reservation utility ru_j , or the largest possible perceived utility improvement from Agent i 's offers, $\Delta_j = u_j(x_{s_i(t)}^*) - u_j(x_i^t)$, is larger than $s_j(t) - u_j(x_{ru_j}^*)$. Thus, Agent i is not sure about whether negotiations will result in an agreement if Agent i stops conceding at any period t before reaching the agent's own reservation utility. Because reaching an agreement would provide higher utility than reaching no agreement, by Definition 13, it is weakly incentive compatible for Agent i to continue to concede. Therefore, it is weakly incentive compatible for all the agents to keep conceding throughout the negotiation process. In other words, an agent will not deliberately stop conceding. \square

Our convergence result for the sequential projection strategy for offer generation, as presented in Section 3, is a general result that does not depend on

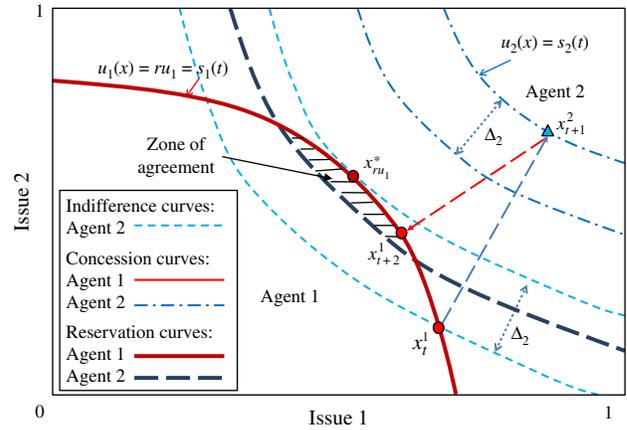


Figure 6 (Color online) The Negotiation Will Not Stall When Agent 1 Concedes to Reservation Utility

specific features of the concession strategy as long as all agents concede to their reservation utilities and the zone of agreement is nonempty. Hence, it also applies to the reactive concession strategy. In other words, as long as all agents using the reactive concession strategy keep conceding to their reservation utilities, an agreement is guaranteed. Theoretically, we have proven that none of the agents will *deliberately* stop conceding, but we have not proven whether the agents may *reactively* stop conceding. By our design of reactive concession strategy, if Agent i 's offer provides a higher utility than Agent j 's reservation utility, Agent j will keep conceding. Additionally, as long as Agent i keeps conceding and Agent j perceives the concession of Agent i , Agent j will still reactively continue to concede. The special case that Agent j may not perceive the concession of Agent i may happen because Agent i concedes to its reservation utility at early stage.

Figure 6 illustrates an example of bilateral negotiation in which the negotiation will not stall even if Agent 1 concedes to the agent's own reservation utility immediately in period t . The offer x_t^1 that Agent 1 proposes in period t lies on Agent 1's reservation indifference curve but not in the zone of agreement. By the nature of the projection offer-generation strategy, the distance between the offers of Agent 1 and Agent 2 decreases over time. The offers that Agent 1 proposes (e.g., offer x_{t+2}^1 in the figure) gradually enter the zone of agreement. By the design of our reactive concession strategy, Agent 2 keeps conceding and an agreement is achieved.

To illustrate the above example is not restricted to bilateral negotiation. Figure 7 illustrates an example of negotiation among three agents. In Figure 7, Agent 1 concedes to the agent's own reservation utility very early in period t . Offer x_t^1 is out of the zone of agreement. Using the sequential projection strategies, the standing offers of the agents become closer

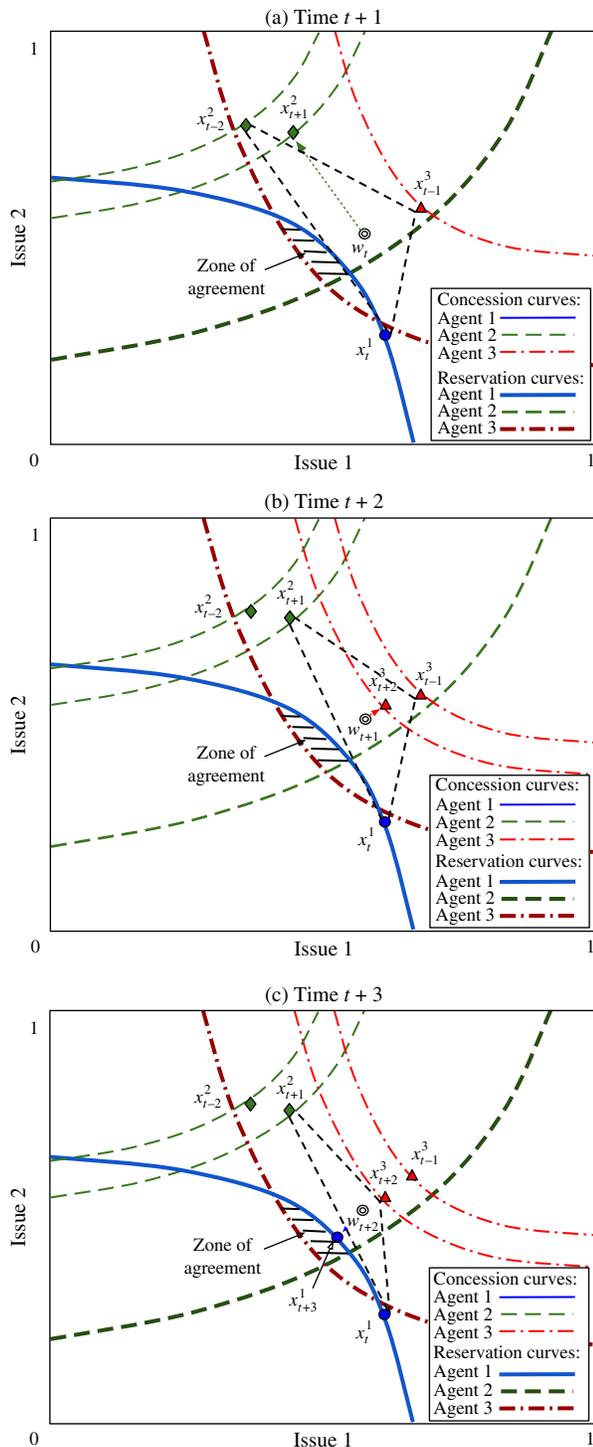


Figure 7 (Color online) Negotiation Will Not Stall for Negotiation Among Three Agents Even If Agent 1 Concedes Immediately to Its Reservation Utility, Since Agent 1's Offers Will Gradually Enter the Zone of Agreement

to one another as time increases and both Agent 2 and Agent 3 keep conceding. The next offer that Agent 1 proposes, x_{t+3}^1 , enters the zone of agreement. Therefore, the negotiation will not stall even if Agent 1 concedes to the agent's own reservation utility very

early in the negotiation at period t . We cannot guarantee that the offers on Agent 1's reservation utility curve will always enter the zone of agreement. If Agent 2's offer does not enter the zone of agreement, other agents may *reactively* stop conceding and the negotiation may stall. For example, in Figure 6, if Agent 1's offer does not enter the zone of agreement, Agent 2's maximum concession from period t is bounded by Δ_2 and the negotiation ends without an agreement.

5. Computational Experiments

In the previous sections, we have proven that if the agents negotiate using the sequential projection strategy, they will reach an agreement. In this section, we present the results of multi-agent computational experiments on randomly generated scenarios to explore various issues of convergence, scalability, and robustness of the sequential projection strategy and the reactive concession strategy.

From the perspective of the overall multi-agent system, the parameters of interest and their values in our experiments are as follows: (a) the number of agents varied from 2 to 9; (b) the number of issues varied from 2 to 5; (c) the values of the agents' reservation utility varied from 0.10 to 0.30 with an increment of 0.05; and (d) the number of simulation runs is 100 per simulation scenario. The agents' utility functions follow a very general class of hyperquadric functions (Hanson 1988): $u_k(x) = 1 - \sum_{i=1}^Q |H_i(x)|^{m_i}$, where x is the n -dimensional proposal vector, $H_i(x) = \sum_{j=1}^N a_{ij}x_j$, $n_i = l_i/m_i$, and $l_i, m_i \in \mathbb{Z}^+$; $f(x)$ is strictly concave if $1 < n_i < \infty$. Hyperquadrics are a general class of functions used in computer graphics (Hanson 1988) and can model a wide range of convex functions. Agent k 's feasible set of offers at period t is the intersection of the unit N -dimensional hypercube $[0, 1]^N$ with $u_k(x) \geq s_k(t)$. Commonly used convex utility functions in the economics literature, such as the Cobb-Douglas functions, can be shown to be special cases of the class of hyperquadric functions. The utility functions of the agents are randomly drawn from the class of hyperquadric functions.

We evaluate our solution with respect to the Nash bargaining solution (Nash 1950), which is Pareto optimal (Nash 1950, Roth 1977, Lensberg 1988). The Nash solution maximizes the agents' joint utility (i.e., the product of the utilities). For the class of (strictly) concave utility functions that we consider, the Nash solution can be obtained by solving a convex optimization problem.

Figure 8 shows a typical sequence of offers with a final agreement generated by five agents negotiating on three issues using the reactive concession strategy.

Table 1 demonstrates the performance of the algorithm when the reactive concession strategy is used.

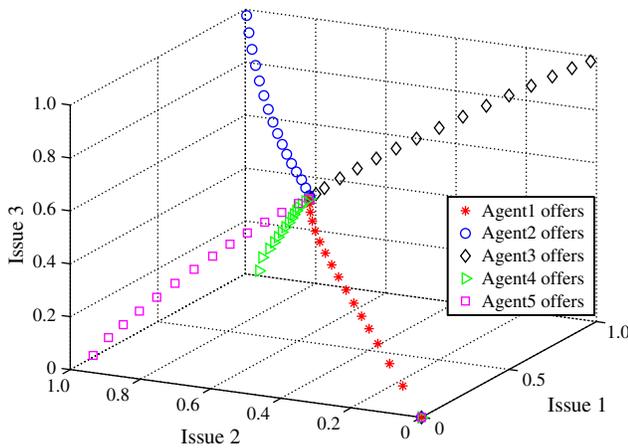


Figure 8 (Color online) Sequence of Offers Made by Five Agents with a Final Agreement in a Three-Issue Negotiation Scenario Using the Sequential Projection Algorithm

Table 1 Performance of Sequential Projection Algorithm with Reactive Concession Strategy

No. of agents	No. of rounds to convergence		Ratio of joint utility	
	Mean	Standard deviation	Mean	Standard deviation
2	62.85	1.99	0.9386	0.0726
3	65.00	1.73	0.9268	0.0307
5	70.43	1.51	0.9098	0.0455
7	74.38	1.87	0.9173	0.0773
9	78.71	1.11	0.9469	0.0169

Here, we varied the number of agents between two and nine, while keeping the number of issues and reservation utilities fixed. The reservation utility of the agents is assumed to be 0.2, and the number of issues is assumed to be 3. The results are averaged over 100 random runs for each row of the table. The numerical tolerance used for convergence is 0.001. As can be seen from Table 1 (second and third columns), the number of rounds required for convergence is fairly stable. The fourth column gives the ratio of our solution to the Nash solution. The quality of the solutions is satisfactorily close to that of the Nash bargaining solutions.

To check the robustness of our findings when varying the number of issues, we performed a sensitivity analysis by varying the number of issues and the reservation utility of the agents. Table 2 shows the performance of the algorithm for five agents negotiating on different numbers of issues, from two to five. The reservation utility of the agents is assumed to be 0.2. The results in Table 2 show that the number of rounds is quite stable if we increase the number of issues. Table 3 shows the performance of the algorithm for five agents negotiating on three issues with reservation utilities varying from 0.10 to 0.30 with an

Table 2 Performance of Sequential Projection Algorithm with Five Agents Negotiating on a Different Number of Issues

No. of issues	No. of rounds to convergence		Ratio of joint utility	
	Mean	Standard deviation	Mean	Standard deviation
2	69.86	1.21	0.9271	0.0391
3	70.43	1.51	0.9098	0.0455
4	71.14	1.68	0.9655	0.0226
5	70.85	1.95	0.9494	0.0222

Table 3 Performance of Sequential Projection Algorithm with Five Agents Negotiating on Three Issues with Different Values of Reservation Utility

Reservation utility	No. of rounds to convergence		Ratio of joint utility	
	Mean	Standard deviation	Mean	Standard deviation
0.10	63.43	1.27	0.9484	0.0318
0.15	66.86	1.07	0.9423	0.0294
0.20	70.43	1.51	0.9098	0.0455
0.25	75.14	1.57	0.9058	0.0640
0.30	83.83	6.46	0.8698	0.1004

increment of 0.05. The results in Table 3 show the number of rounds is quite stable even if we vary the reservation utilities of the agents, as long as the zone of agreement is nonempty.

To substantiate our claim that the sequential projection strategy will converge not contingent on the concession strategy, we performed experiments in which the agents concede in a random fashion. Table 4 shows the performance of the sequential projection algorithm with five agents negotiating on three issues with one or more agents (shown in the first column) conceding in a random fashion to their reservation utilities. The reservation utility of the agents is assumed to be 0.2. Specifically, we let those agents draw at random a desirable utility between the original desirable utility and their reservation value. The simulation results confirm the convergence of our proposed algorithm in the presence of randomness in the concession strategy.

To check the robustness of the sequential projection strategy to noise in the calculation of the projection, we let one or more agents add a random noise term ε to their algorithmically generated bids. In generating the random noise, we use the uniform norm of ε , $\|\varepsilon\|_\infty$, to be smaller than 0.01. We also restrict the Euclidean norm of ε , $\|\varepsilon\|$, to be smaller than the distance between the projection x_t^i and w_{t-1} for Agent i in period t . Table 5 shows the performance of the sequential projection algorithm with five agents negotiating on three issues where one to three agents use

Table 4 Performance of Sequential Projection Algorithm with Five Agents Negotiating on Three Issues with One or More Agents Conceding in a Random Fashion to Reservation Utility

No. of agents conceding in a random fashion	No. of rounds to convergence		Ratio of joint utility	
	Mean	Standard deviation	Mean	Standard deviation
	1	72.43	4.16	0.9315
2	70.29	2.14	0.9566	0.0343
3	65.71	3.20	0.9554	0.0318

Table 5 Performance of Sequential Projection Algorithm with Five Agents Negotiating on Three Issues when One or More Agents Use Noisy Bids

No. of agents using noisy bids	No. of rounds to convergence		Ratio of joint utility	
	Mean	Standard deviation	Mean	Standard deviation
	1	69.14	2.41	0.9340
2	70.43	1.90	0.9115	0.0529
3	70.86	1.77	0.9351	0.0243

Table 6 Performance of Reactive Concession Strategy with Five Agents Negotiating on Three Issues when One or More Agents Concede Immediately to Their Reservation Utilities

No. of agents conceding immediately to reservation values	No. of rounds to convergence		Ratio of joint utility	
	Mean	Standard deviation	Mean	Standard deviation
	1	65.86	7.80	0.9332
2	67.15	2.88	0.8808	0.0566
3	66.29	4.39	0.8713	0.0885

noisy bids. The reservation utility of the agents is assumed to be 0.2. The simulation results confirm the convergence of our proposed algorithm in the presence of noisy bids.

The goal of the final set of experiments was to test the robustness of the reactive concession strategy. In particular, if one agent arrives at reservation utility very early, compared with the other agents, the first agent will not be able to make nonzero concessions out of necessity, not deliberately. Table 6 shows the performance of the sequential projection algorithm with five agents negotiating on three issues, where one to three agents converge immediately to their reservation values. Each agent's reservation utility is 0.2. Although we cannot provably guarantee stalling will not happen, we show that in our many simulation scenarios, stalling never arose. To further explore the robustness of the reactive concession strategy, we simulated negotiation on three issues with a different number of agents and different values of reservation

Table 7 Percentage of Convergence of Reactive Concession Strategy for Three-Issue Negotiations when One Agent Converges Immediately to Its Reservation Utility

Reservation utilities	No. of agents (%)	
	3	5
0.10	100	100
0.15	100	100
0.20	100	100
0.25	100	96
0.30	100	92

utility when one agent concedes immediately to reservation utility. The results in Table 7 show the negotiation scenarios never stall for negotiation among three agents. Only a very small percentage of negotiation scenarios, namely 4% and 8%, respectively, stall only for the negotiations among five agents with corresponding reservation utility of 0.25 and 0.30.

6. Conclusions and Future Direction

In this paper, we advance the computational literature on multilateral negotiation in several aspects. First, we propose a distributed negotiation strategy for general multilateral, multi-attribute negotiation where agents have no knowledge about the other players' utility functions. In particular, we propose a sequential projection strategy for offer generation, and prove that by using this strategy, the agents are *guaranteed* to arrive at a satisficing agreement, if the agents concede to their reservation utilities and if the zone of agreement is nonempty. This convergence property is guaranteed irrespective of the concession strategies the agents use, under concave utility functions, and despite the fact that agents have no information about the preferences of other agents. We argue that the sequential projection strategy is rational for agents to follow. In considering agents' incentives to concede, we propose and analyze a reactive concession strategy and prove that none of the agents has incentive to deliberately stop conceding. We have also performed computational experiments and demonstrated that, in practice, in randomly generated problem instances, the quality of the solution that our algorithm obtains is quite close to the Nash bargaining solution (that maximizes the joint utility of the agents). The negotiation converges well in a reasonable number of iterations and scales as the number of agents or number of issues is increased. Furthermore, we tested the robustness of our algorithms to various deviation strategies, namely, agents conceding randomly, agents generating noisy bids, and agents conceding immediately to their reservation utilities.

Methodologically, we advance the alternating projection algorithms literature in that we consider *multiple* sets that *move*; we are the first to prove the convergence property under this novel setting.

Our work can be extended in several directions, one of which is to explore the design of a postnegotiation phase for finding a solution that is better for all the agents after a satisficing agreement has been reached. A second direction is to design negotiation strategies in the presence of different best alternatives to a negotiated agreement. A third direction is to design rational strategies for agents to negotiate in the presence of hard deadlines. For negotiation with no information in the presence of deadlines, there are simple examples that show that any concession strategy that allows an agent to reach the agent's reservation utility by the deadline cannot guarantee that the agents will reach an agreement. Thus, this poses a challenge in identifying a class of provably good negotiation strategies for negotiating in the presence of hard deadlines. A fourth direction is to extend the sequential projection strategy to negotiation between multiple negotiation teams. Lastly, it may be useful to incorporate learning; specifically, agents could try to learn other agents' preferences during negotiation to increase negotiation efficiency.

Acknowledgments

This paper extends our previous work (Zheng et al. 2013b) on two sets in bilateral multi-issue negotiation. A succinct version of this work without the proofs and details appeared as an extended abstract in Zheng et al. (2013a). The authors thank Maqbool Dada, and conference participants of the 2015 Group Decision and Negotiation International Conference at the Warsaw School of Economics for their constructive comments. The authors also thank Ram Ramesh, the associate editor, and three anonymous reviewers for their high-quality reviews that helped improved the paper. This research was made possible by funding from The Army Research Institute ARO [Grant W911NF-08-1-0301].

Appendix

For the purpose of completeness, we present the convex optimization formulation for finding the Nash bargaining solution. Nash (1950) provides an axiomatic approach to define reasonable outcomes in a negotiation, which has since been advanced in many subsequent works. In this paper, we use a convex optimization approach for computing the Nash bargaining solution. Hence, we will restrict our discussion to the formulation of the optimization problem. Let m agents be negotiating on n issues, with the issues taking on continuous values between 0 and 1. Let $u_i(x)$ be Agent i 's utility function, which is assumed to be concave. Without loss of generality, we assume that no agreement results in a utility of 0. The objective function to be maximized is the joint utility, namely, $f(x) = \prod_{i=1}^m u_i(x)$. Because $u_i(x)$ is concave and nonnegative, $f(x)$ is nonnegative, and hence maximizing $f(x)$ is equivalent to maximizing $\log(f(x))$. Let x_j denote the j th component of x . The

convex optimization problem to be solved for computing the Nash equilibrium is

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^m \log(u_i(x)) \\ & \text{s.t.} && u_i(x) \geq ru_i \quad i = 1, \dots, m, \\ & && 0 \leq x_j \leq 1, \quad \forall j = 1, \dots, n, \end{aligned} \quad (25)$$

where ru_i is Agent i 's ultimate reservation utility. Because each u_i is a concave function of x , the set of constraints in (25) forms a convex set. The objective function to be maximized is a sum of log-concave functions, and hence the problem is a convex optimization problem. In the paper, we have used the solver CVX (Grant and Boyd 2015) implemented in MATLAB to obtain the Nash bargaining solutions.

References

- Aknine S, Pinson S, Shakun MF (2004) An extended multi-agent negotiation protocol. *Autonomous Agents Multi-Agent Systems* 8(1):5–45.
- Bauschke H, Borwein J (1996) On projection algorithms for solving convex feasibility problems. *SIAM Rev.* 38(3):367–426.
- Binmore K (1985) Bargaining and coalitions. Roth AE, ed. *Game-Theoretic Models of Bargaining* (Cambridge University Press, Cambridge, UK), 269–304.
- Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press, Cambridge, UK).
- Buffett S, Spencer B (2005) Learning opponents' preferences in multi-object automated negotiation. *Seventh Internat. Conf. Electronic Commerce (ICEC'05), Xi'an, China*.
- Chalamish M, Kraus S (2012) Automed: An automated mediator for multi-issue bilateral negotiations. *Autonomous Agents Multi-Agent Systems* 24(3):536–564.
- Chari K, Agrawal M (2007) Multi-issue automated negotiations using agents. *INFORMS J. Comput.* 19(4):588–595.
- Chatterjee K, Samuelson W (1983) Bargaining under incomplete information. *Oper. Res.* 31(5):835–851.
- Chen L, Pu P (2004) Survey of preference elicitation methods. Technical report IC/200467, Swiss Federal Institute of Technology, Lausanne.
- Cheney W, Goldstein A (1959) Proximity maps for convex sets. *Proc. Amer. Math. Soc.* 10(3):448–450.
- Combettes PL (1997) Hilbertian convex feasibility problem: Convergence of projection methods. *Appl. Math. Optim.* 35(3):311–330.
- Endriss U (2006) Monotonic concession protocols for multilateral negotiation. *Proc. Fifth Internat. Joint Conf. Autonomous Agents Multiagent Systems, AAMAS '06, Hakodate, Hokkaido, Japan*, 392–399.
- Fatima S, Wooldridge MJ, Jennings NR (2004) An agenda-based framework for multi-issue negotiation. *Artif. Intell.* 152(1):1–45.
- Fatima S, Wooldridge M, Jennings N (2006) Multi-issue negotiation with deadlines. *J. Artif. Intell. Res.* 27(1):381–417.
- Grant M, Boyd S (2015) CVX: MATLAB software for disciplined convex programming, version 2.1. Accessed March 31, 2016, <http://cvxr.com/cvx>.
- Hanson AJ (1988) Hyperquadrics: Smoothly deformable shapes with convex polyhedral bounds. *Comput. Vision, Graphics, Image Processing* 44(2):191–210.
- Harsanyi JC, Selten R (1972) A generalized Nash solution for two-person bargaining games with incomplete information. *Management Sci.* 18(5, Part 2):80–106.
- Heiskanen P, Ehtamo H, Hamalainen RP (2001) Constraint proposal method for computing Pareto solutions in multi-party negotiations. *Eur. J. Oper. Res.* 133(1):44–61.

- Ito T, Hattori H, Klein M (2007) Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces. *Proc. 20th Internat. Joint Conf. Artificial Intelligence (IJCAI'07), Hyderabad, India*, 1347–1352.
- Kalai E (1977) Proportional solutions to bargaining situations: Inter-temporal utility comparisons. *Econometrica* 45(7):1623–1630.
- Kalai E, Smorodinsky M (1975) Other solutions to Nash's bargaining problem. *Econometrica* 43(3):513–518.
- Klein M, Faratin P, Sayama H, Bar-Yam Y (2003) Negotiating complex contracts. *Group Decision Negotiation* 12(2):111–125.
- Lai G, Sycara K (2009) A generic framework for automated multi-attribute negotiation. *Group Decision Negotiation* 18(2):169–187.
- Lange A, Löschel A, Vogt C, Ziegler A (2010) On the self-interested use of equity in international climate negotiations. *Eur. Econom. Rev.* 54(3):359–375.
- Lensberg T (1988) Stability and the Nash solution. *J. Econom. Theory* 45(2):330–341.
- Li C, Tesauro G (2003) A strategic decision model for multi-attribute bilateral negotiation with alternating offers. Nisan N, ed. *Proc. ACM Conf. Electronic Commerce* (ACM Press, New York), 208–209.
- Lin R, Kraus S, Wilkenfeld J, Barry J (2008) Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Art. Intell.* 172(6–7): 823–851.
- Lopes de Lima M, Camponogara E, Limon Marruedo D, Munoz de la Pena D (2015) Distributed satisficing MPC. *Control Systems Tech., IEEE Trans.* 23(1):305–312.
- Myerson RB (1979) Incentive compatibility and the bargaining problem. *Econometrica* 47(1):61–73.
- Nash J (1950) The bargaining problem. *Econometrica* 18(2):155–162.
- Roth AE (1977) Individual rationality and Nash's solution to the bargaining problem. *Math. Oper. Res.* 2(1):64–65.
- Rubinstein A (1982) Perfect equilibrium in a bargaining model. *Econometrica* 50(1):97–109.
- Shakun MF (2005) Multi-bilateral multi-issue e-negotiation in e-commerce with a tit-for-tat computer agent. *Group Decision Negotiation* 14(5):383–392.
- Simon HA (1956) Rational choice and the structure of the environment. *Psych. Rev.* 63(2):129–138.
- Stirling WC (2003) *Satisficing Games and Decision Making: With Applications to Engineering and Computer Science* (Cambridge University Press, Cambridge, UK).
- Stirling WC (2005) Social utility functions-part i: Theory. *Systems, Man, Cybernetics, Part C: Appl. Rev., IEEE Trans.* 35(4):522–532.
- Sycara K, Dai T (2010) Agent reasoning in negotiation. Kilgour D, Eden C, eds. *Handbook of Group Decision and Negotiation* (Springer, New York), 437–451.
- Wu M, de Weerd MM, Poutre HL (2009) Efficient methods for multi-agent multi-issue negotiation: Allocating resources. Yang JJ, Yokoo M, Ito T, Jin Z, Scerri P, eds. *Proc. PRIMA, Lecture Notes in Computer Science*, Vol. 5925 (Springer, New York), 97–112.
- Zeng DD (2001) Managing flexibility for inter-organizational electronic commerce. *Electronic Commerce Res.* 1(1–2):33–51.
- Zheng R, Chakraborty N, Dai T, Sycara K (2013a) Multiagent negotiation on multiple issues with incomplete information. *Proc. Twelfth Internat. Conf. Autonomous Agents Multiagent Systems, Saint Paul, MN*, 1279–1280.
- Zheng R, Chakraborty N, Dai T, Sycara K, Lewis M (2013b) Automated bilateral multi-issue negotiation with no information about opponent. *Proc. Hawaii Internat. Conf. Systems Sci., Wailea, Maui, HI*, 520–527.