

## MIT Open Access Articles

*A Secure and Fair Protocol that Addresses Weaknesses of the Nash Bargaining Solution in Nonlinear Negotiation*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Fujita, Katsuhide, Takayuki Ito, and Mark Klein. "A Secure and Fair Protocol That Addresses Weaknesses of the Nash Bargaining Solution in Nonlinear Negotiation." *Group Decis Negot* 21, no. 1 (March 3, 2010): 29–47.

**As Published:** <http://dx.doi.org/10.1007/s10726-010-9194-6>

**Publisher:** Springer Netherlands

**Persistent URL:** <http://hdl.handle.net/1721.1/105777>

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

**Terms of Use:** Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



# A Secure and Fair Protocol that Addresses Weaknesses of the Nash Bargaining Solution in Nonlinear Negotiation

Katsuhide Fujita · Takayuki Ito · Mark Klein

Published online: 3 March 2010  
© Springer Science+Business Media B.V. 2010

**Abstract** Negotiation with multiple interdependent issues is an important problem since much of real-world negotiation falls into this category. This paper examines the problem that, in such domains, agent utility functions are nonlinear, and thereby can create nonconvex Pareto frontiers. This in turn implies that the Nash Bargaining Solution, which has been viewed as the gold standard for identifying a unique optimal negotiation outcome, does *not* serve that role in nonlinear domains. In nonlinear domains, unlike linear ones, there can be multiple Nash Bargaining Solutions, and all can be sub-optimal with respect to social welfare and fairness. In this paper, we propose a novel negotiation protocol called SFMP (the Secure and Fair Mediator Protocol) that addresses this challenge, enabling secure multilateral negotiations with fair and pareto-optimal outcomes in nonlinear domains. The protocol works by (1) using nonlinear optimization, combined with a Multi-Party protocol, to find the Pareto front without revealing agent's private utility information, and (2) selecting the agreement from the Pareto set that maximizes a fair division criterion we call approximated fairness. We demonstrate that SFMP is able to find agreements that maximize fairness

---

Takayuki Ito—Visiting from Nagoya Institute of Technology, Japan.

---

K. Fujita (✉)  
School of Techno-Business Administration, Nagoya Institute of Technology, Gokiso, Showa-ku,  
Nagoya, Aichi 466-8555, Japan  
e-mail: fujita@itolab.mta.nitech.ac.jp

T. Ito · M. Klein  
Center for Collective Intelligence, MIT Sloan School of Management, 5 Cambridge Center,  
NE25-749A, Cambridge, MA 02445, USA  
e-mail: ito.takayuki@nitech.ac.jp; takayuki@mit.edu

M. Klein  
e-mail: m\_klein@mit.edu

and social welfare in nonlinear domains, and out-performs (in terms of outcomes and scalability) previously developed nonlinear negotiation protocols.

**Keywords** Multi-issue negotiation · Nonlinear utility function · Fairness

## 1 Introduction

Negotiation is an important aspect of daily life and represents an important topic in the field of multi-agent systems research. There has been extensive work in the area of automated negotiation; that is, where automated agents negotiate with other agents in such contexts as e-commerce (Kraus 2001), large scale argumentation (Malone et al. 2007), collaborative design, and so on. Even though many contributions have been made in this area (Bosse and Jonker 2005; Faratin et al. 2002; Fatima et al. 2004; Soh and Li 2004) most have dealt exclusively with negotiations involving one or more independent issues with simple (linear) utility functions. Many real-world negotiations, however, are more complex, involving interdependent issues and therefore, nonlinear utility functions. When designers work together to design a car, for example, the utility of a given carburetor choice is highly dependent on which engine is chosen. The key impact of such issue dependencies is that they result in agent utility functions that are nonlinear, i.e., that have multiple optima. Most existing negotiation protocols, though well-suited for linear utility functions, work poorly when applied to nonlinear problems (Klein et al. 2003).

The Nash bargaining solution, which maximizes the product of the agent utilities, is a well-known metric that provably identifies the optimal (fair and social-welfare-maximizing) agreement for negotiations in linear domains (Kaneko 1980; Binmore et al. 1986; Nash 1950). In *nonlinear* domains, however, the Pareto frontier will often not satisfy the convexity assumption required to make the Nash solution optimal and unique (Denicolo and Mariotti 2000; Kaneko 1980; Nash 1950). There can, in other words, be multiple agreements in nonlinear domains that satisfy the Nash Bargaining Solution, and many or all of these will have sub-optimal fairness and/or social welfare. We need, therefore, a new approach if we want to produce good outcomes for nonlinear negotiations.

In this paper, we present a secure mediated protocol (the Secure and Fair Mediator Protocol, or SFMP) that addresses this challenge. The protocol consists of two main steps. In the first step, SFMP uses a nonlinear optimizer, integrated with a secure information sharing technique called the Multi-Party Protocol (Shamir 1979), to find the Pareto front without causing agents to reveal private utility information. In the second step, an agreement is selected from the set of Pareto-optimal contracts using a metric, which we call approximate fairness, that measures how equally the total utility is divided across the negotiating agents (Robertson 1998 etc.). We demonstrate that SFMP produces better scalability and social welfare values than previous nonlinear negotiation protocols.

The remainder of this paper is organized as follows. First, we describe a model of nonlinear negotiation with utility functions based on constraints, and show how the Nash Bargaining Solution can lead to sub-optimal results in such contexts. Second,

we describe a new protocol (SFMP) designed to address this challenge. Finally, we present the experimental results, describe related work, and draw conclusions.

## 2 Negotiation with Nonlinear Utility Functions

We consider the situation where  $n$  agents want to reach an agreement with a mediator who manages the negotiation from a man-in-the-middle position. There are  $M$  issues  $S = \{s_1, \dots, s_M\}$  to be negotiated. The number of issues represents the number of dimensions in the utility space. The issues are shared: all agents are potentially interested in the values for all  $M$  issues. Each issue  $s_j$  has a value drawn from the domain of integers  $[0, X]$ , i.e.,  $s_j \in \{0, 1, \dots, X\} (1 \leq j \leq M)$ .<sup>1</sup> A contract is represented by a vector of issue values  $\mathbf{s} = (s_1, \dots, s_M)$ . We assume that agents have an incentive to cooperate to achieve win-win agreements because a non-agreement has lower utility than an agreement.


An agent's utility function, in our formulation, is described in terms of constraints. There are  $l$  constraints,  $c_k \in C$ . Each constraint represents a region in the contract space with one or more dimensions and an associated utility value. Constraint  $c_k$  has value  $w_i(c_k, \mathbf{s})$  if and only if it is satisfied by contract  $\mathbf{s} (1 \leq k \leq l)$ . Every agent has its own, typically unique, set of constraints.

An agent's utility for contract  $\mathbf{s}$  is defined as the sum of the utility for all the constraints it satisfies, i.e., as  $u_i(\mathbf{s}) = \sum_{c_k \in C, \mathbf{s} \in x(c_k)} w_i(c_k, \mathbf{s})$ , where  $x(c_k)$  is a set of possible contracts (solutions) of  $c_k$ . This expression produces a "bumpy" nonlinear utility function with high points where many constraints are satisfied and lower regions where few or no constraints are satisfied. This represents a crucial departure from previous efforts on multi-issue negotiation, where contract utility is calculated as the weighted sum of the utilities for individual issues, producing utility functions shaped like flat hyper planes with a single optimum.

Figure 1 shows an example of a utility function generated via a collection of binary constraints involving Issues 1 and 2. Constraint A, for example, which has a value of 55, holds if the value for Issue 1 is in the range [3,7] and the value for Issue 2 is in the range [4,6]. The utility function is highly nonlinear with many hills and valleys. For our work, we assume that many real-world utility functions are more complex than this, involving more than two issues as well as higher-order (e.g., trinary and quaternary) constraints. In recent work (e.g., Marsa-Maestre et al. 2009a), several types of constraints were proposed.

This constraint-based utility function representation allows us to capture the issue interdependencies common in real world negotiations. Constraint A, for example, captures the fact that a value of 4 is desirable for issue 1 *if* issue 2 has the value 4, 5 or 6. Note, however, that this representation is also capable of capturing linear utility functions as a special case (they can be captured as a series of unary constraints).

<sup>1</sup> A discrete domain can come arbitrarily close to a real domain by increasing its size. As a practical matter, many real-world issues that are theoretically 'real' numbers (delivery date, cost) are discretized during negotiations.



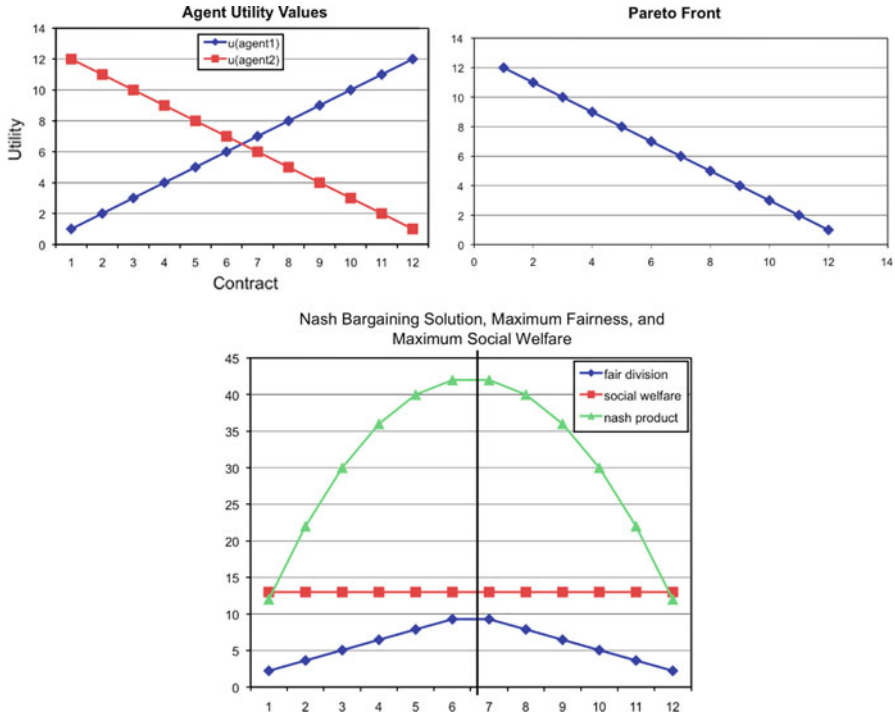
**Fig. 1** An example of a nonlinear utility function generated using constraints

A negotiation protocol for complex contracts can, therefore, also handle simple (linear) contract negotiations.

Working in the nonlinear domain has a number of important impacts on the kind of negotiation protocols that can be effective. First, consider pareto-optimality. Pareto-optimality is widely recognized as a basic requirement for a good negotiation outcome. It is defined as follows: Contract  $\mathbf{s} = (s_1, \dots, s_M)$  is **Pareto optimal** if there is no  $\mathbf{s}'$  such that  $u_i(\mathbf{s}') > u_i(\mathbf{s})$  for all agents ( $u_i(\mathbf{s})$  is agent  $i$ 's utility value). Pareto-optimality thus eliminates all contracts where there are others that are better for all the parties involved. In a linear negotiation (i.e., where the agent utility functions are defined as the weighted sum of the values for each issue), it is computationally trivial to find the Pareto front, and the social welfare (sum of agent utilities) for every contract on the Pareto frontier is the same. In fact, the Pareto-optimal frontier for a negotiation will be “sparse” in our model, i.e., the Pareto-optimal contracts points will be few in number and widely scattered.

Next, let us consider fairness. Fairness is critical in bargaining theory because some experimental results suggest that it deeply influences human decision-making (Ken et al. 75 etc.) in such contexts as family decision making (e.g., where will we go on our next vacation?), the less formal economy of consumer transactions (such as ticket scalpers or flea markets), and price setting for consumer purchases. The ultimatum game is a popular example of this effect (Alvard 2004; Bolton 1991). People tend to offer “fair” (i.e., 50:50) splits, and offers less than 20% are often rejected in this game, even though it is irrational in this game to reject any deal, since the alternative is a zero payoff. There are many other studies about the relationship between decision making and “fairness” in the experimental economics and behavioral economics fields (Keeney and Raiffa 1993; Werner et al. 1982).

The Nash Bargaining Solution (i.e., the contract that maximizes the “Nash product” = the product of the agent’s utility functions) is a widely-used approach

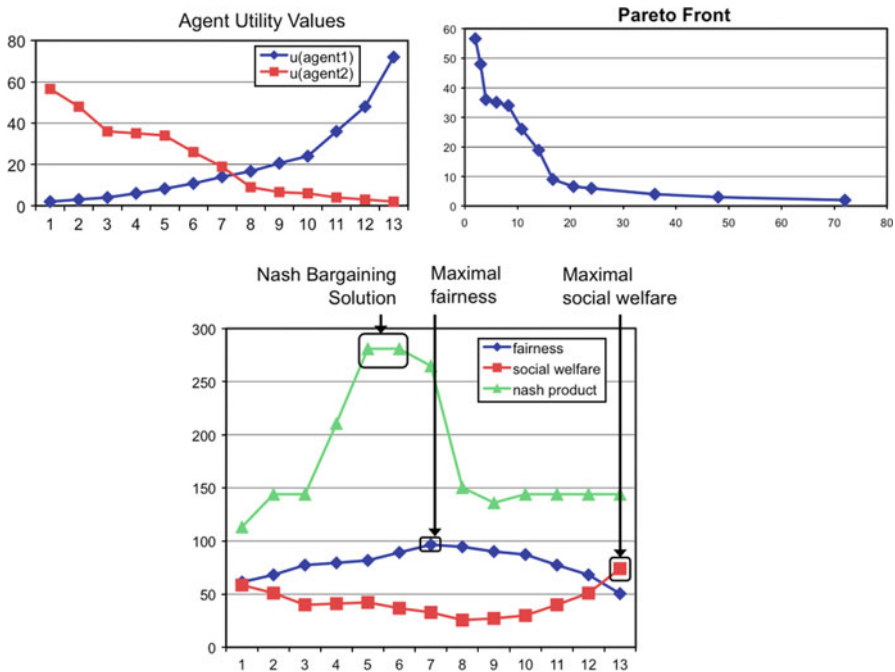


**Fig. 2** The relationship of nash product, fairness and social welfare in a linear utility function

for identifying the most “fair” contract from those that make up the Pareto front. As we can see in Fig. 2, the Nash Bargaining Solution divides utility equally amongst the negotiating parties, in a linear domain. It can be proven, in fact, that there is a unique Nash bargaining solution for negotiations with convex pareto fronts, which is satisfied trivially for negotiations with linear utilities (Nash 1950).<sup>2</sup>

These properties change radically in nonlinear negotiation. As we can see in Fig. 3, when agents have nonlinear utility functions, the Pareto front can be non-convex (Myerson 1991). There can be multiple Nash bargaining solutions, even with continuous issue domains, and some of the Nash bargaining solutions may be non-optimal in terms of social welfare and fair division of utility. It is even straightforward to find nonlinear cases where *all* the contracts on the Pareto front are Nash bargaining solutions, despite the fact that many of them diverge widely from maximal fairness and social welfare. The Nash Bargaining Solution concept, which is widely used as a basis for negotiation protocols for linear domains, will thus often fare poorly in nonlinear domains. We need, therefore, to find negotiation protocols that can achieve high social welfare and fairness values with nonlinear agent utilities.

<sup>2</sup> In negotiations with discretized issue domains, there can be multiple Nash Bargaining Solutions, but they will all be clustered right next to each other and thus offer very similar fairness values.



**Fig. 3** The relationship of nash product, fairness and social welfare in a nonlinear utility function

### 3 Secure and Fair Mediator Protocol with Approximated Fairness

The Secure and Fair Mediator Protocol (SFMP) was defined to achieve these goals while also protecting agent's private utility information. SFMP consists of two main steps: (1) finding the set of Pareto-optimal contracts, and (2) selecting a fair contract from that set. These steps are defined below.

#### 3.1 Finding the Pareto Front

This step is achieved using a mediated approach [Fujita et al. \(2008a, 2009\)](#). One or more mediators propose contracts, initially randomly generated, and ask the agents which ones they prefer. The mediators use this preference information to provide the objective function for a nonlinear optimization technique such as simulated annealing or a genetic algorithm. Over the course of multiple rounds, the mediators converge on the set of pareto-optimal contracts. We assume, as is common in negotiation contexts, that agents prefer not to share their utility functions with others, in order to preserve a competitive edge. Accordingly, our protocol uses a Multi-Party Protocol [Shamir \(1979\)](#) to ensure that mediators can calculate the sum of the agents' utilities without learning, or revealing, the individual agents' utility information. A detailed explanation of the Multi-Party Protocol is given in "Appendix 1".

### 3.2 Selecting the Final Agreement

SFMP selects the final agreement from the Pareto-optimal contract set by calculating which one is the most fair. Several definitions of “fair” have been identified in social choice and game theory [Robertson \(1998\)](#). Suppose we have a division  $X = X_1 \cup \dots \cup X_n$  among  $n$  agents where agent  $i$  receives  $X_i$ . “Simple” fair division results if  $u_i(X_i) \geq 1/n$  whenever  $1 \leq i \leq n$  (each agent gets at least  $1/n$ .) Another definition, from game theory, calls a division  $X$  is fair if and only if it is Pareto-optimal and envy-free ([Chevalleyre et al. 2007](#)). A division is “envy-free” if no agent feels another has a strictly larger piece of the utility ([Robertson 1998](#)).

We adopt simple fair division as our concept of fairness. Contract agreements, in general, rarely fully satisfy this condition. We measure, accordingly, how *close* an agreement is to simple fair division by calculating its “approximated fairness”, i.e., the deviation of each agent’s utility from the average of the total utility. The approximate fairness of a contract is defined, formally, as follows:

$$V(u_1, \dots, u_n) = \sum_{i=1}^n \frac{(u_i - \bar{u})^2}{n}$$

( $u_1, \dots, u_n$ : agent’s utility value in contract,  $\bar{u}$ : the average of all agent’s utility value).

An ideal contract, therefore, has an approximated fairness value of zero, and all other contracts will have larger values. The final agreement selected by our protocol is the contract from the Pareto-optimal set with the smallest approximated fairness value.

Note that our fairness concept is equivalent to the Nash bargaining solution in linear contexts with continuous issue domains. Assume that  $u_1 + u_2 + \dots + u_n = K$  (constant) (where  $u_i$ : agent  $i$ ’s utility value). The Nash product is maximized when  $u_1 = u_2 = \dots = u_n = K/n$  (this has been proven mathematically in the field of Iso-perimetric Problems). The key difference is that our measure generalizes to nonlinear domains. Approximated fairness does not, however, correspond to the Kalai-Smorodinsky solution because the latter isn’t always fair ([Thomson 1992](#)).

## 4 Experiments

We ran a series of negotiation simulation experiments in order to demonstrate the weaknesses of the Nash Bargaining Solution in nonlinear domains, and to compare the performance of the SFMP protocol we defined against that of previous approaches. The subsections below describe the experiment setup and results.

### 4.1 Detailed Description of Secure and Fair Mediator Protocol (SFMP)

The SFMP protocol utilizes multiple mediators in order to help assure agent privacy. We assume that there are  $k = mn$  mediators  $M_j$  and  $n$  agents ( $A_i$ ), where  $m$  is an



arbitrary integer. Note that this approach requires that  $m$  is relatively high if we wish to effectively conceal the agent's private information. If the number of mediators is low, the chances increase that all of the mediators will collude, and thus compromise the agent's privacy.

**(Optional Pre-Negotiation Step) Contract space division among mediators:** The mediators divide the contract space between them, so each mediator searches a different subregion. Suppose, for example, that there are two issues whose domain is the integers from 0 to 10. In this case, mediator 1 can manage the region of values 0–5 for issue 1 and values 0–10 for issue 2, while mediator 2 can manage the region of values 6–10 for issue 1 and values 0–10 for issue 2. This step is optional, but it has the advantage of potentially reducing the time needed to search the contract space by allowing parallel computation.

**(Step 1) Secure search to find a Pareto-optimal contract set:** Each mediator searches its assigned portion of the contract space using a local search algorithm (Russell and Norvig 2002). We employed Hill Climbing (HC), Simulated Annealing (SA), and Genetic Algorithm (GA) in our experiments. In HC, an agent starts with a random solution and, at each step, makes some random mutations and selects the one that causes the greatest utility increase. When the algorithm cannot find any more improvements, it terminates. In SA, each step of the SA algorithm replaces the current solution by a randomly generated nearby contract, with a probability that depends on the change in utility value and a global parameter  $T$  (the virtual temperature) that is gradually decreased during the process. The agent moves almost randomly when the temperature is high, but acts increasingly like a hill climber as the temperature decreases. When  $T$  is 0, the search is terminated. The advantage of SA is that it is able to avoid getting stuck in the local optima that occur in nonlinear optimization problems, and often finds more optimal solutions than hill climbing. GA is a search technique inspired by evolutionary biology, using such techniques as inheritance, mutation, selection, and crossover. Initially many individual contracts are randomly generated to form an initial population. After that, at each step, a proportion of the existing population is selected, based on their 'fitness' (i.e., utility values). Crossover and mutation is then applied to these selections to generate the next generation of contracts. This process is repeated until a termination condition has been reached. The objective function for all these local search algorithms is the maximization of social welfare. At each search step, the mediators determine the social welfare values by securely gather the utility values for the current contract(s) from their assigned agents. We call this **secure value gathering**.

**(Step 2) Identify agreement:** All mediators share the maximum value in their sub-region of the contract space with all the other mediators. Based on that, they identify the pareto-optimal contract set. The mediators then select the contract, in that set, that minimizes the approximated fairness metric. This represents the final agreement for that negotiation.

## 4.2 Nash Product Maximization Search

For a comparison case, we used Nash Product Maximization Search (NPMS) to find the Nash bargaining solution for our tests (Russell and Norvig 2002). Our implementation used simulated annealing to maximize the Nash product for the negotiating agents, gathering their utility values using the secure multi-party protocol. Simulated annealing has been shown to be very effective for nonlinear optimization tasks (Ito et al. 2007). We can use the results of NPMS to assess the scale of the performance decrement caused by using the Nash Bargaining Solution concept in nonlinear domains.

## 4.3 Settings

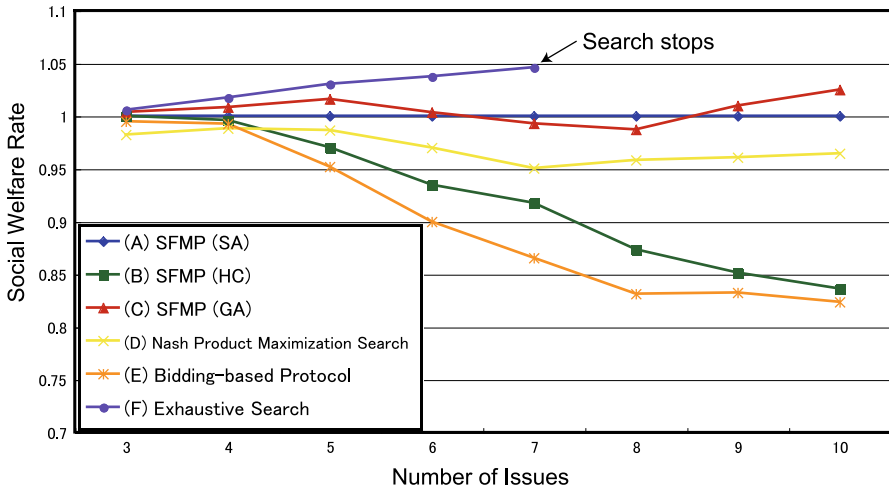
We conducted five experiments to evaluate the effectiveness of our approach. In each experiment, we ran 100 negotiations between agents with randomly generated utility functions. The number of agents was six, and the number of mediators was four. The mediators could calculate the sum of the agent's utility if four mediators got together. The search space was divided equally amongst the mediators. The domain for the issue values was  $[0,9]$ . The constraints included 10 unary constraints, 5 binary constraints, 5 trinary constraints, and so on (a unary constraint relates to one issue, a binary constraint relates to two issues, and so on). The maximum value for a constraint was  $100 \times (\text{number of issues})$ . Constraints that satisfy many issues thus have, on average, larger utility, which seems reasonable for many domains. In a meeting scheduling domain, for example, higher order constraints concern more people than lower order constraints, so they are more important. The maximum width for a constraint was 7. The following constraints, for example, are all valid: Issue 1 =  $[2,6]$ , Issue 3 =  $[2, 9]$ .

We compared the following negotiation protocols: SFMP (SA), SFMP (HC), SFMP (GA), Nash Product Maximization Search (NPMS), Bidding-based protocol, and Exhaustive Search.

**(A) SFMP (SA):** "SFMP (SA)" is SFMP using Simulated Annealing as the optimization algorithm. The initial temperature was 50 degrees. The temperature decreased 0.1 degrees for each iteration, resulting in 500 iterations.  $20 + (\text{Number of issues}) \times 5$  searches were conducted, with the initial start point changed randomly for each search.

**(B) SFMP (HC):** "SFMP (HC)" is SFMP using Hill Climbing as the optimization algorithm. We employed the random restart hill climbing mechanism (Russell and Norvig 2002).  $20 + (\text{number of issues}) \times 5$  searches were conducted, with the initial start point changed randomly for each search.

**(C) SFMP (GA):** "SFMP (GA)" is SFMP using a Genetic Algorithm as the optimization algorithm. The population size was  $20 + (\text{number of issues}) \times 5$ . We employed a basic crossover method in which two parent individuals were combined to produce two children (one-point crossover). The fitness function was the sum of all agents' (declared) utility. 500 iterations were conducted. Mutations happened at very small probability. In a mutation, one of the issues in a contract vector was randomly chosen and changed.



**Fig. 4** Comparison of social welfare

**(D) Nash Product Maximization Search (NPMS):** “Nash Product Maximization Search” used SA to search for the Nash bargaining solution(s), i.e., for contracts that maximize the Nash product. The initial temperature was 50 degrees. For each iteration, the temperature decreased 0.1 degree, resulting in 500 iterations.  $20 + (\text{Number of issues}) \times 5$  searches were conducted, with the initial start point changed randomly for each search. These settings are the same as those for SFMP (SA).

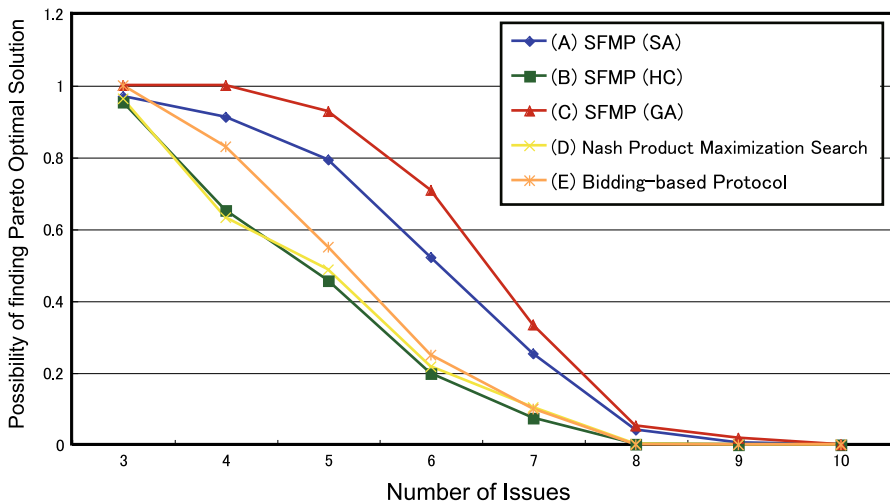
**(E) Bidding-based protocol:** “Bidding-based protocol” is the protocol proposed in Ito et al. (2007). In this protocol, the number of samples taken during random sampling is  $(\text{number of issues}) \times 200$ . The threshold used to remove contract points that have low utility is 200. The limitation on the number of bids per agent is  $\sqrt[3]{6,400,000}$  for  $N$  agents. This method fails to reach agreements if the number of issues is more than eight because this method has too much computational complexity.

**(F) Exhaustive Search:** “Exhaustive search” is a centralized brute force algorithm that traverse the entire contract search space to find the Pareto-optimal contract set. The final contract was then selected using our approximated fairness measure. This approach was only computationally practical when the number of issues was seven or fewer.

Our code was implemented in Java 2 (1.5) and run on a core 2-duo processor iMac with 1.0 GB memory on the Mac OS  $\times$  10.5 operating system.

#### 4.4 Experimental Results

Figure 4 compares the social welfare achieved by these six methods. The evaluation measure we used was the (social welfare for final agreement from method)/(social welfare for final agreement from SFMP (SA)). As predicted, we found that SFMP (SA) and SFMP (GA) performed better than NPMS, confirming our claim that the

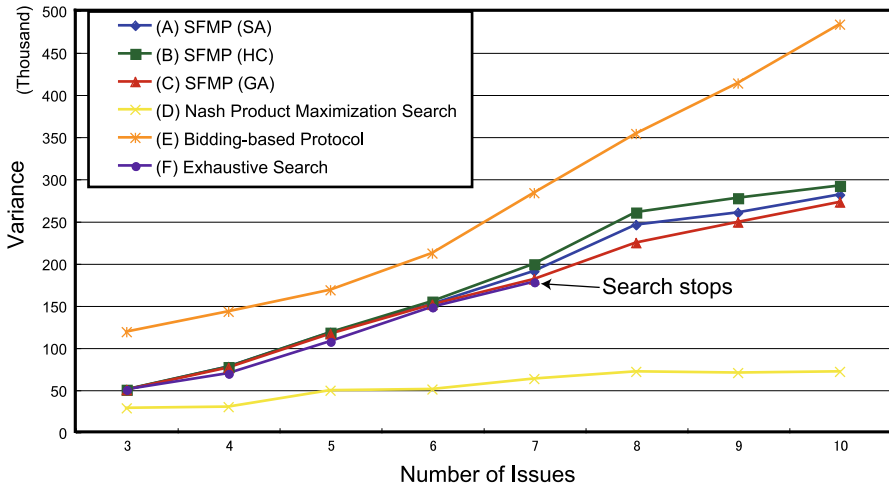


**Fig. 5** Finding Pareto-optimal contracts

Nash Bargaining Solution produces sub-optimal outcomes when applied to nonlinear negotiation. SFMP (SA) and SFMP (GA) had about equal performance. Neither produced fully optimal results, reflecting the difficulty of performing optimization in large nonlinear contract spaces. All the SFMP protocols performed better than the Basic Bidding Protocol, which was hampered by the limit on the number of bids per agent necessitated by the combinatorics of winner determination in this protocol. The performance of SFMP (HC) decreased rapidly as the number of issues grew, because hill climbing got stuck on local optima. The performance of SFMP (SA) and SFMP (GA) did not decrease appreciably as the number of issues increased.

Figure 5 compares the number of Pareto-optimal contracts found by the six methods. In this experiment, we limited the domain of each issue to  $[0,4]$ , so all Pareto-optimal contracts could be found, in a reasonable amount of time, using the exhaustive search algorithm. We found that SFMP(SA) and SFMP(GA) were better at finding Pareto-optimal contracts than either the Nash Product Maximization Search or the Bidding-Based Protocol. This makes sense, since the SFMP was explicitly designed to find the entire Pareto front first, before selecting a final agreement, while the other protocols were not. We also found that SFMP(SA) and SFMP(GA) performed better than the bidding-based protocol, because the latter often fails to find Pareto-optimal solutions due to the limit on the number of bids allowed by each agent. As always, the performance of SFMP (HC) decreased rapidly as the number of issues grew. SFMP (GA) showed the highest performance on this measure, because GA is inherently more suitable for finding Pareto-optimal contract sets. However, for all methods, when the number of issues increased, the percentage of Pareto-optimal contracts found decreased drastically.

Figure 6 assessed fairness by comparing the variance of the agent's utilities for the final agreements, across the six methods. Lower variance is better, because it means that utility is distributed more fairly across the agents. The SFMP protocols showed better performance than the "Bidding-based Protocol" on this measure because the basic



**Fig. 6** Comparison of variance

bidding protocol doesn't consider fairness when finding agreements. SFMP (GA) showed the lowest (best) value among the SFMP variants. NPMS outperformed the SFMPs on this measure. This is counter to what we predicted: in nonlinear domains, we would expect the Nash bargaining solutions to vary widely in their fairness values, causing NPMS to produce, on the average, sub-optimal fairness values.

We can potentially explain these results by considering the allocation of computational effort in nonlinear optimization. In an even moderately large nonlinear optimization problem, the contract space is too large to explore exhaustively. If we have only 10 issues with 10 possible values per issue, for example, this produces a space of  $10^{10}$  (10 billion) possible contracts. As a result, with limited computational resources, we have no guarantee of finding the complete Pareto front. SFMP is presumably only able to find a subset of the Pareto-optimal contracts, and those are scattered over the entire frontier. Because the coverage is sparse, SFMP will often not happen to find the Pareto-optimal contract that optimizes the fairness metric. This will reduce the average fairness score for SFMP. NPMS, by contrast, devotes its entire computational effort to finding a single Nash-product-maximizing contract. Even though it is an inferior optimization objective, it has the benefit of a more concentrated application of computing resources.

This interpretation is supported by Fig. 7, which shows the utility values for SFMP and NPMS for a case with two agents and five issues, with randomly generated nonlinear utility functions. The red points show the contracts considered by NPMS, while the blue points show the contracts considered by SFMP. Since SFMP aims to find the entire Pareto front, it searches throughout the Pareto frontier. NPMS, by contrast, aims to find the contract that directly maximizes the Nash product, so it focuses its search toward the middle of the Pareto frontier. As Fig. 7 shows, SFMP in this case got closer to the Pareto frontier than NPMS.

Figure 8 compares the failure rates across the six methods, to assess their scalability of our methods. For all the methods, if the computing time method exceeded 100 s,

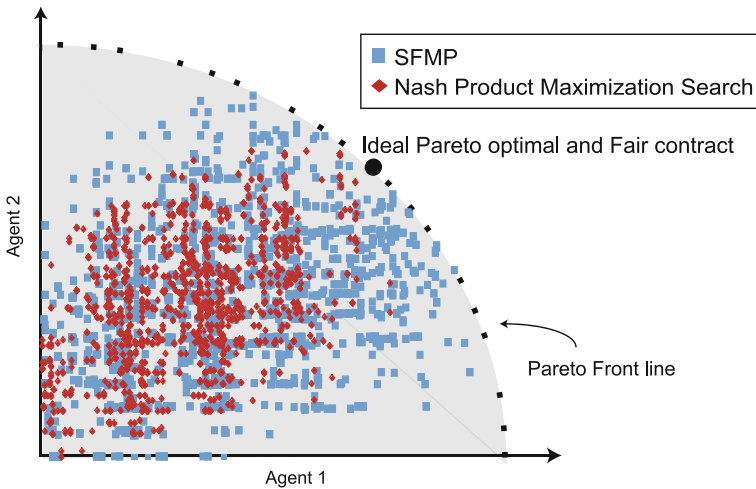


Fig. 7 Comparison between SFMP and Nash product maximization search

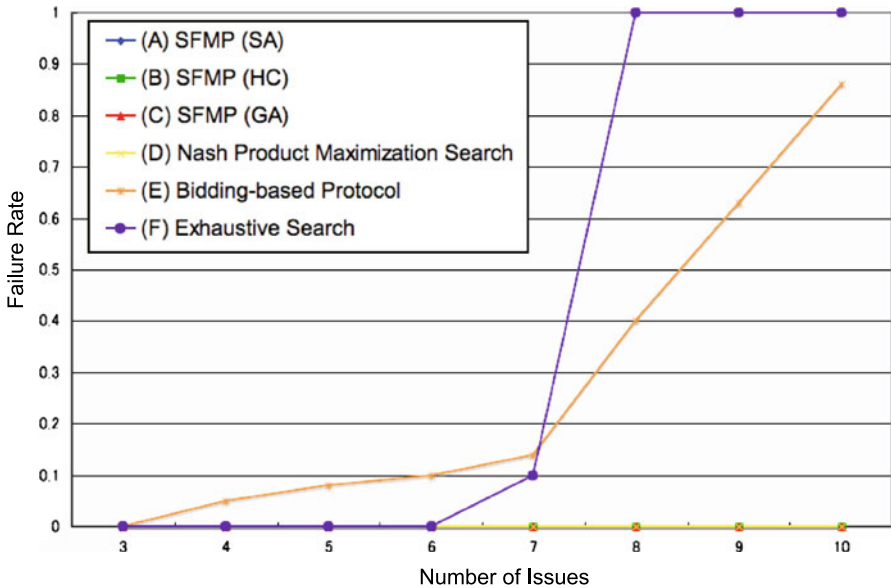


Fig. 8 Comparison of failure rate

the negotiation was aborted and it was treated as a failure. The failure rate for the bidding-based and exhaustive search protocols increased exponentially with the number of issues. This is because that the computational complexity of finding agreements in these protocols is quite large. All the other protocols had negligible failure rates.

The key experimental results can be summarized as follows:

- SFMP, as predicted, maximizes social welfare more effectively than NPMS. It also out-performs the bidding-based protocol.

- “SFMP” finds fairer contracts than the “Bidding-based protocol”, but is less fair than NPMS.
- “SFMP” has a lower failure rate (and thus greater scalability) than the “Bidding-based Protocol”.

We also found that the negotiation methods were sensitive to the complexity of negotiation setting, due to the combinatorics of the local search algorithms they employed. The larger the number of issues, the lower the optimality of the outcomes.

## 5 Related Work

There has been some recent work on agent privacy in distributed constraint optimization problems (DCOP) field (Greenstadt et al. 2006; Maheswaran et al. 2005). Even though negotiation seems to involve a straightforward constraint optimization problem, we have been unable to exploit existing work on high efficiency constraint optimizers. Such solvers attempt to find the solutions that maximize the weights of the satisfied constraints, but they do not account for the fact that agents are all self-interested.

Klein et al. (2003) proposed a mediated protocol that uses simulated annealing to achieve near-optimal outcomes with binary issue interdependencies in large contract spaces. This work was limited, however, to bilateral negotiations.

Ito et al. (2007) proposed a bidding-based protocol for negotiation with multiple interdependent issues. In this protocol, agents generate bids by sampling their utility functions, and the mediator finds the optimum combination of submitted bids from agents. Marsa-Maestre et al. (2009a, b) proposed a mechanism that balances utility and deal probability for the bidding and deal identification processes. This mechanism improves the solution optimality and scalability compared with bidding-based protocols. Fujita et al. (2008b) proposed a protocol where the mediator selects representatives who propose alternatives to other agents. This protocol increases scalability by reducing in effect the number of agents who have to come to an agreement. These protocols, however, focus on finding high social welfare contracts and don't consider fairness and Nash-bargaining solution features.

Ponsati and Watson (1997) proposed four procedures (global bargaining, separate bargaining, simultaneous implementation, and independent implementation) by which the bargaining may take place in two-person, multiple-issue bargaining problems. Kraus and Schechter (2003) considered a model of bilateral negotiation, but in a complete information setting. Their work focuses mostly on examining agent strategies when one agent loses utility over time, while the other gains utility over time. These papers all focus, in addition, on bilateral multi-issue negotiations., while the work reported in this paper focuses on negotiations among more than two agents.

Lin and Chou (2003) explored a range of protocols based on mutation and selection over binary contracts. This paper does not describe what kind of utility function is used, nor does it present any experimental analyses, so it remains unclear whether this strategy enables sufficient exploration of utility space. Barbuceanu and Lo (2000) presented an approach based on constraint relaxation. However, this paper provides no experimental analysis and merely presents a small toy problem with 27 contracts. Debenham (2004) proposed a multi-issue bargaining strategy that models iterative

information gathering that takes place during negotiation. However, these models are not explicitly designed to address the problem of complex and high dimensional negotiations. Klein et al. (2003) presented a protocol applied with near optimal results to medium-sized bilateral negotiations with binary dependencies. This work demonstrated both scalability and high optimality values for multilateral negotiations and higher order dependencies.

Lai et al. (2006a, b) presented a protocol for multi-issue problems for bilateral negotiations. Robu et al. (2005), Robu and Poutre (2006) presented a multi-item and multi-issue negotiation protocol for bilateral negotiations in electronic commerce situations. Gerding et al. (2006) proposed a negotiation mechanism where the bargaining strategy is decomposed into a concession strategy and a Pareto-search strategy. However, these papers also focus on bilateral multi-issue negotiations.

Shew and Larson (2008) proposed multi-issue negotiation that employs a third-party as a mediator to guide agents toward equitable solutions. This framework also employs an agenda that serves as a schedule for the ordering of issue negotiation. Agendas are very interesting because agents only need to focus on a few issues. This paper also focuses on bilateral negotiations, however, this framework can apply to the negotiations among more than two agents.

Jonker et al. (2007) proposed a negotiation model called ABMP that can be characterized as cooperative one-to-one multi-criteria negotiation in which the privacy of both parties is protected as much as desired. Hindriks et al. (2006) proposed an approach based on a weighted approximation technique to simplify the utility space. The resulting approximated utility function without dependencies can be handled by negotiation algorithms that can efficiently deal with independent multiple issues and have a polynomial time complexity (Jonker and Treur 2001). Hindriks et al. (2008) proposed a checking procedure to mitigate this risk and showed that by tuning this procedure's parameters, outcome deviation can be controlled. These studies reflect interesting viewpoints, but they only focused on bilateral trading or negotiations. These approaches are efficient, however, the utility function of this approaches is different from the one based on multi-dimensional constraints.

Fatima et al. (2007a, b) proposed bilateral multi-issue negotiations with time constraints. This method can find approximate equilibrium in polynomial time where the utility function is nonlinear. However, these papers focused on bilateral multi-issue negotiations. Li et al. (2009) proposed a method in which the mediator searches for a compromise direction based on an Equal Directional Derivative approach and computes a new tentative agreement in bilateral multi-issue negotiations. However, this method only focused on multilateral negotiation.

## 6 Conclusions

We show, in this paper, that the Nash Bargaining Solution, although provably optimal for negotiations with linear utilities, can lead to sub-optimal outcomes when applied to nonlinear negotiations. We also present the Secure and Fair Mediator Protocol (SFMP), a novel negotiation protocol that utilizes a combination of nonlinear optimization, secure information sharing, and an approximated fairness metric, and demon-



strate that it achieves higher social welfare values than a protocol based on searching for the Nash bargaining solution. Finally, we demonstrate that SFMP out-performs our own previous efforts to enable multi-lateral negotiations in complex domains.

Future work includes building protocols that can find Pareto-optimal contracts more quickly, making them more scalable, and increasing the fairness performance. One potential approach to this problem is to focus the search efforts of the mediators more closely on the fair portion of the Pareto frontier. We plan to investigate incentive-compatibility issues in more detail, to ensure that the protocol can not be “gamed” by agents seeking to gain disproportionate influence or to sabotage the outcomes. Finally, we plan to explore the consequences of the fact that nonlinear problems, unlike linear ones, can produce situations where you have to decide if social welfare or fairness is more important. We will explore protocols that can deal with this situation somehow, for example for giving negotiators the Pareto front and letting them bargain using traditional iterative concession techniques.

### Appendix 1: Secure Value Gathering

Appendix 1 below includes an explanation of secure value gathering. Figure 9 shows an example with three agents and two mediators ( $k = 2$ ). This is just for illustrative purposes: in practical situations,  $k$  should be larger to reduce the likelihood of mediator collusion. In the following,  $u_i$  is agent  $i$ 's utility value.

1. The mediators ask the agents to generate “shares”  $v$ . Each agent  $A_i$  will send one share  $v_{i,j}$  to each mediator  $M_j$ .

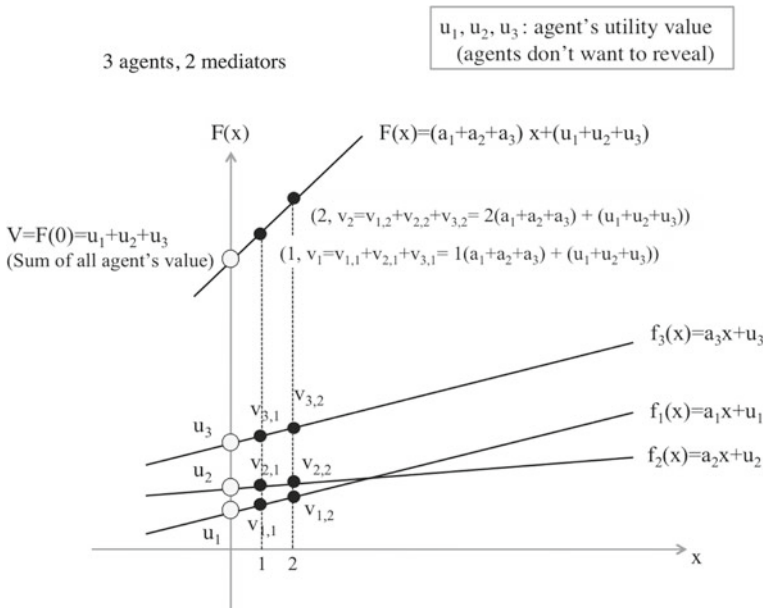


Fig. 9 The example of secure value gathering

2. Each agent  $i(A_i)$  randomly selects a  $k$  dimensional polynomial formula which fulfills  $f_i(0) = u_i$ . In Fig. 9, for example, agent 1 selected  $f_1(x) = a_1x + u_1$ , agent 2 selected  $f_2(x) = a_2x + u_2$  and agent 3 selected  $f_3(x) = a_3x + u_3$ .
3. Each agent ( $A_i$ ) calculates a share  $v_{i,j} = f_i(j)$  for each mediator ( $M_j$ ) and sends it to that mediator. For example, agent  $A_1$ 's share for mediator  $M_2$  would be  $v_{1,2} = f_1(2) = 2a_1 + u_1$ .
4. Every mediator  $j(M_j)$  sums the shares  $v_{1,j}, \dots, v_{n,j}$  it receives from the agents in order to calculate  $v_j = v_{1,j} + \dots + v_{n,j}$ . In Fig. 9, for example, mediator 1 received the shares  $v_{1,1}$ ,  $v_{2,1}$ , and  $v_{3,1}$  and calculated  $v_1 = v_{1,1} + v_{2,1} + v_{3,1}$ .
5. The  $j$  mediators add their share sums  $v_j$  together to calculate  $F(x)$  for  $x$  from 1 to  $j$ . Using Lagrange's interpolating polynomial, it is then straightforward to calculate  $F(0)$ , which corresponds to the sum of all the agent's utility values for a contract. The net result is that the social welfare is calculated without any one mediator knowing the utility of any contract for any individual agent.

## References

- Alvard MS (2004) 14. the ultimatum game, fairness, and cooperation among big game hunters. *Found Hum Sociality* 1(9):413–451
- Barbuceanu M, Lo WK (2000) Multi-attribute utility theoretic negotiation for electronic commerce. In: *Proceedings of the AAMAS workshop on agent mediated electronic commerce (AMEC-2000)*, pp 15–30
- Binmore K, Rubinstein A, Wolinsky A (1986) The nash bargaining solution in economic modelling. *The RAND Corp J RAND J Econ* 17:176–188
- Bolton GE (1991) A comparative model of bargaining: theory and evidence. *Am Econ Rev* 81(5):1096–1136
- Bosse T, Jonker CM (2005) Human vs. computer behaviour in multi-issue negotiation. In: *Proceedings of 1st international workshop on rational, robust, and secure negotiations in multi-agent systems (RRS-2005)*, pp 11–24
- Chevalyre Y, Endriss U, Estivie S, Maudet N (2007) Reaching envy-free states in distributed negotiation settings. In: *Proceedings of 20th international joint conference on artificial intelligence (IJCAI-2007)*, pp 1239–1244
- Debenham J (2004) Bargaining with information. In: *Proceedings of the 3rd international joint conference on autonomous agents and multi-agent systems (AAMAS-2004)*, pp 663–670
- Denicolo V, Mariotti M (2000) Nash bargaining theory, nonconvex problems and social welfare orderings. *Theory Decis* 48(4):351–358
- Faratin P, Sierra C, Jennings NR (2002) Using similarity criteria to make issue trade-offs in automated negotiations. *Artif Intell* 142:205–237
- Fatima SS, Wooldridge M, Jennings NR (2004) Optimal negotiation of multiple issues in incomplete information settings. In: *Proceedings of third international joint conference on autonomous agent and multi-agent systems (AAMAS-2004)*, pp 1080–1087
- Fatima SS, Wooldridge M, Jennings NR (2007a) An analysis of feasible solutions for multi-issue negotiation involving nonlinear utility functions. In: *Proceedings of the 8th international joint conference on autonomous agents and multi-agent systems (AAMAS-2009)*, pp 1041–1048
- Fatima SS, Wooldridge M, Jennings NR (2007b) Approximate and online multi-issue negotiation. In: *Proceedings of the 6th international joint conference on autonomous agents and multi-agent systems (AAMAS-2007)*, pp 947–954 (2007)
- Fujita K, Ito T, Klein M (2008a) Preliminary result on secure protocols for multiple issue negotiation problems. In: *Proceedings of the 11th Pacific rim international conference on multi-agents (PRIMA-2008)*
- Fujita K, Ito T, Klein M (2008b) A representative-based multi-round protocol for multi-issue negotiations. In: *Proceedings of the 7th international joint conference on autonomous agents and multi-agent systems (AAMAS-2008)*, pp 1573–1576

- Fujita K, Ito T, Klein M (2009) Approximately fair and secure protocols for multiple interdependent issues negotiation. In: Proceedings of the eighth international joint conference on autonomous agents and multi-agent systems (AAMAS-2009), pp 1287–1288
- Gerding E, Somefun D, Poutre HL (2006) Efficient methods for automated multi-issue negotiation: negotiating over a two-part tariff. *Int J Intell Syst* 21:99–119
- Greenstadt R, Pearce J, Tambe M (2006) Analysis of privacy loss in distributed constraint optimization. In: Proceedings of The 21th association for the advancement of artificial intelligence (AAAI-2006)
- Hindriks K, Jonker C, Tykhonov D (2006) Interdependencies between issues for multi-issue negotiation. *Cooperative information agents X, Lect Notes Comput Sci* 4149:301–316
- Hindriks K, Jonker C, Tykhonov D (2008) Avoiding approximation errors in multi-issue negotiation with issue dependencies. In: Proceedings of the 1st international workshop on agent-based complex automated negotiations (ACAN-2008)
- Ito T, Hattori H, Klein M (2007) Multi-issue negotiation protocol for agents: exploring nonlinear utility spaces. In: Proceedings of 20th international joint conference on artificial intelligence (IJCAI-2007), pp 1347–1352
- Jonker CM, Robu V, Treur J (2007) An agent architecture for multi-attribute negotiation using incomplete preference information. *J Auton Agent Multi Agent Syst* 15:221–252
- Jonker CM, Treur J (2001) An agent architecture for multi-attribute negotiation. In: Proceedings of 17th international joint conference on artificial intelligence (IJCAI-2001), pp 1195–1201
- Kaneko M (1980) An extension of the nash bargaining problem and the nash social welfare function. *Theory Decis* 12(2):135–148
- Keeney RL, Raiffa H (1993) *Decisions with multiple objectives*. Cambridge University Press, Cambridge
- Ken B, Shaked A, Sutton J (1975) Testing noncooperative bargaining theory: a preliminary study. *Am Econ Rev* 75:1178–1180
- Klein M, Faratin P, Sayama H, Bar-Yam Y (2003) Negotiating complex contracts. *Group Decis Negot* 12(2):58–73
- Kraus S (2001) *Strategic negotiation in multiagent environments*. Cambridge University Press, Cambridge
- Kraus S, Schechter O (2003) Strategic-negotiation for sharing a resource between two agents. *Comput Intell* 19:9–41
- Lai G, Li C, Sycara K (2006a) A general model for pareto optimal multi-attribute negotiations. In: Proceedings of the 2nd international workshop on rational, robust, and secure negotiations in multi-agent systems (RRS-2006)
- Lai G, Sycara K, Li C (2006b) A decentralized model for multi-attribute negotiations with incomplete information and general utility functions. In: Proceedings of the 2nd international workshop on rational, robust, and secure negotiations in multi-agent systems (RRS-2006)
- Li M, Vo QB, Kowalczyk R (2009) Searching for fair joint gains in agent-based negotiation. In: Proceedings of the 8th international joint conference on autonomous agents and multi-agent systems (AAMAS-2009), pp 1049–1056
- Lin RJ, Chou ST (2003) Bilateral multi-issue negotiations in a dynamic environment. In: Proceedings of the AAMAS workshop on agent mediated electronic commerce (AMEC-2003)
- Maheswaran RT, Pearce JP, Varakantham P, Bowring E (2005) Valuations of possible states (vps): a quantitative framework for analysis of privacy loss among collaborative personal assistant agents. In: Proceedings of the 4th international joint conference on autonomous agents and multi-agent systems (AAMAS-2005)
- Malone TW, Klein M (2007) Harnessing collective intelligence to address global climate change. *Innov J* 2(3):15–26
- Marsa-Maestre I, nad Juan R, Velasco MALC, de la Hoz E (2009a) Effective bidding and deal identification for negotiations in highly nonlinear scenarios. In: Proceedings of the 8th international joint conference on autonomous agents and multi-agent systems (AAMAS-2009), pp 1057–1064
- Marsa-Maestre I, Lopez-Carmona MA, Velasco JR, Ito T, Fujita K, Klein M (2009b) Balancing utility and deal probability for negotiations in highly nonlinear utility spaces. In: Proceedings of twenty-first international joint conference on artificial intelligence (IJCAI-09)
- Myerson R (1991) *Game theory: analysis of conflict*. Harvard University Press, Harvard
- Nash J (1950) The bargaining problem. *Econometrica* 18:155–162
- Ponsati C, Watson J (1997) Multiple-issue bargaining and axiomatic solutions. *Int J Game Theory* 26:501–524
- Robertson J, Webb W (1998) *Cake-cutting algorithms: be fair if you san*. A K Peters Ltd, Natick

- Robu V, Poutre HL (2006) Retrieving the structure of utility graphs used in multi-item negotiation through collaborative filtering of aggregate buyer preferences. In: Proceedings of the 2nd international workshop on rational, robust, and secure negotiations in multi-agent systems (RRS-2006)
- Robu V, Somefun DJA, Poutre JL (2005) Modeling complex multi-issue negotiations using utility graphs. In: Proceedings of the 4th international joint conference on autonomous agents and multi-agent systems (AAMAS 2005)
- Russell SJ, Norvig P (2002) Artificial intelligence: a modern approach. Prentice Hall, New Jersey
- Shamir A (1979) How to share a secret. *Commun ACM* 22(11):612–613
- Shew J, Larson K (2008) The blind leading the blind: a third-party model for bilateral multi-issue negotiations under incomplete information. In: Proceedings of the 1st international workshop on agent-based complex automated negotiations (ACAN-2008)
- Soh LK, Li X (2004) Adaptive, confidence-based multiagent negotiation strategy. In: Proceedings of third international joint conference on autonomous agent and multi-agent systems (AAMAS-2004), pp 1048–1055
- Thomson W (1992) Cooperative models of bargaining. *Handb Game Theory Econ Appl* 2:1238–1284
- Werner G, Rolf S, Bernd S (1982) An experimental analysis of ultimatum bargaining. *J Econ Behav Organ* 3:367–388