Avoiding the Prisoner's Dilemma in Auction-based Negotiations for Highly Rugged Utility Spaces

Ivan Marsa-Maestre, Miguel A. Lopez-Carmona, Juan R. Velasco and Enrique de la Hoz
Computer Engineering Department, Universidad de Alcala
Edificio Politecnico, Ctra. N-II, Km 36.500
28871 Alcala de Henares (Madrid), Spain
{ivan.marsa,miguelangel.lopez,juanramon.velasco,enrique.delahoz}@uah.es

ABSTRACT

There is a number of recent research lines addressing complex negotiations in highly rugged utility spaces. However, most of them focus on overcoming the problems imposed by the complexity of the scenario, without analyzing the strategic behavior of the agents in the models they propose. Analyzing the dynamics of the negotiation process when agents with different strategies interact is necessary to apply these models to real, competitive environments, where agents cannot be supposed to behave in the same way. Specially problematic are situations like the well-known prisoner's dilemma, or more generally, situations of high price of anarchy. These situations imply that individual rationality drives the agents towards strategies which yield low individual and social welfares. In highly rugged scenarios, such situations usually make agents fail to reach an agreement, and therefore negotiation mechanisms should be designed to avoid them. This paper performs a strategy analysis of an auction-based negotiation model designed for highly rugged scenarios, revealing that the approach is prone to the prisoner's dilemma. In addition, a set of techniques to solve this problem are proposed, and an experimental evaluation is performed to validate the adequacy of the proposed approaches to improve the strategic stability of the negotiation process.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search—*heuristic methods*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*multiagent systems*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*coherence and coordination*

General Terms

Algorithms, Design, Experimentation

Keywords

multi-agent systems, multi-issue negotiation, highly-nonlinear utility spaces

Cite as: Avoiding the Prisoner's Dilemma in Auction-based Negotiations for Highly Rugged Utility Spaces, Ivan Marsa-Maestre *et al.*, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 425-432

1. INTRODUCTION

In the last years, there has been an increasing interest in complex negotiations which involve multiple negotiating parties and multiple, interdependent issues [5]. Specially challenging are those scenarios involving highly rugged utility spaces, since traditional negotiation approaches (mostly intended for linear or quasi-concave utility functions) cannot be applied to these more complex scenarios.

We can find some successful research works in the literature addressing negotiation in complex utility spaces. In [4], an auction-based protocol is proposed for nonlinear utility spaces generated using weighted constraints. This approach is based on taking random samples of the contract space and applying simulated annealing to these samples to identify high utility regions for each agent, sending these regions as bids to a mediator, and then performing a search in the mediator to find overlaps between the bids of the different agents. In a similar scenario [8], we proposed to take samples from the constraints space instead. Experiments show that these approaches achieve high effectiveness (measured as high optimality rates and low failure rates for the negotiations) in moderately rugged utility spaces.

In [9], we joined efforts with the the aforementioned authors to address highly-rugged utility spaces. We proposed the use of a *quality factor* to balance utility and deal probability in the negotiation process. This quality factor is used to bias bid generation and deal identification taking into account the agents' attitudes towards risk (i.e. allowing agents to give more importance to utility or to deal probability depending on their own attitudes towards risks). The experiments show that this balance between utility and deal probability greatly improves the effectiveness of the negotiation in highly-rugged utility spaces.

However, the proposed approach draws several concerns. Though the quality factor is supposed to be able to model agents' risk attitudes, the experiments limit these attitudes to a somewhat "cooperative" environment, where all agents have the same, neutral risk attitude. In a real, competitive environment, we expect to have agents with different risk attitudes interacting. This raises the problem of agent strategic behavior. What happens when risk averse agents interact with risk willing agents? Is there a dominant strategy? If so, does this dominant strategy lead to satisfying solutions, or is the approach prone to the prisoner's dilemma? Furthermore, since the complexity (i.e. ruggedness) of the utility spaces of the agents may also vary, it seems logical to think that agent strategies should vary accordingly. In this paper, we intend to address these questions in the following ways:

- We perform a strategy analysis of the auction-based protocol for constraint-based utility spaces. This analysis allows us to determine the individual dominant strategy and the optimal social strategy for different utility space ruggedness levels. From the results of the analysis we conclude that the auction-based protocol, as described in [9], has stability problems, being prone to the prisoner's dilemma (Section 3).

- We propose a set of mechanisms intended to avoid the prisoner's dilemma in the analyzed protocol. These approaches are based on decoupling the agent's strategies from the deal identification process, by applying different techniques on the mediator after the agents have sent their bids (Section 4).

An experimental evaluation has been performed to validate our hypothesis and evaluate the effects of our contributions. The experimental setting is described in Sections 3.2 and 4.2, along with the discussion of the results obtained. Finally, our proposal is briefly compared to the most closely-related works in the state-of-the-art (Section 5). The last section summarizes our conclusions and sheds light on some future research.

2. AUCTION-BASED NEGOTIATIONS IN HIGHLY RUGGED UTILITY SPACES

The work we propose is a contribution to the strategic behavior of the agents in auction based negotiations for complex utility spaces. In this section, we outline the most relevant previous works our research is related to.

2.1 Constraint-based Nonlinear Utility Spaces

Nonlinear agent preferences can be described by using different categories of functions, like K-additive utility functions, bidding languages, or weighted constraints [9]. In this work we focus on nonlinear utility spaces generated by means of weighted constraints. In these cases, agents' utility functions are described by defining a set of constraints. Each constraint represents a region with one or more dimensions, and has an associated utility value. The number of dimensions of the space is given by the number of issues n under negotiation, and the number of dimensions of each constraint must be lesser or equal than n. The utility yielded by a given potential solution (contract) in the utility space for an agent is the sum of the utility values of all the constraints that are satisfied by that contract. Figure 1 shows a very simple example for two issues and three constraints: a unary constraint $C1$ and two binary constraint $C2$ and $C3$. The utility values associated to the constraints are also shown in the figure. In this example, contract x would yield a utility value for the agent $u(x) = 15$, since it satisfies both $C1$ and $C2$, while contract y would yield a utility value $u(y) = 5$, because it only satisfies $C1$. It can also be noted that unary constraint $C1$ can be seen as a binary constraint where the width of the constraint for issue 2 is all the domain of the issue, so we can generalize and say that all constraints have n dimensions.

More formally, we can define the issues under negotiation as a finite set of variables $x = \{x_i | i = 1, ..., n\}$, and a contract (or a possible solution to the negotiation problem) as

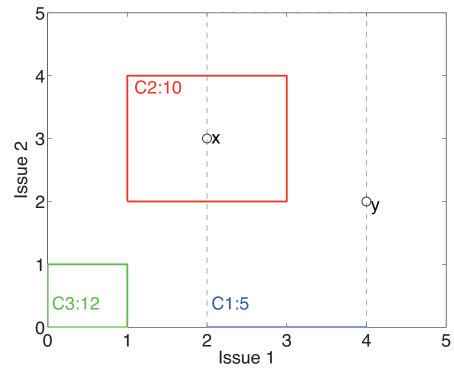

Figure 1: Example of a utility space with two issues and three constraints

a vector $s = \{x_i^s | i = 1, ..., n\}$ defined by the issues' values. Issues take values from the domain of integers $[0, X]$.

Agent utility space is defined as a set of constraints $C = \{c_k | k = 1, ..., l\}$. Each constraint is given by a set of intervals which define the region where a contract must be contained to satisfy the constraint. In this way a constraint c is defined as $c = \{I_i^c | i = 1, ..., n\}$, where $I_i^c = [x_i^{min}, x_i^{max}]$ defines the minimum and maximum values for each issue to satisfy the constraint. Constraints defined in this way describe hyper-rectangular regions in the n-dimensional space. Each constraint c_k has an associated utility value $u(c_k)$.

A contract s satisfies a constraint c if and only if $x_i^s \in I_i^c \forall i$. For notation simplicity, we denote this as $s \in x(c_k)$, meaning that s is in the set of contracts that satisfy c_k. An agent's utility for a contract s is defined as $u(s) = \sum_{c_k \in C | s \in x(c_k)} u(c_k)$, that is, the sum of the utility values of all constraints satisfied by s. This kind of utility functions produces nonlinear utility spaces, with high points where many constraints are satisfied, and lower regions where few or no constraints are satisfied.

2.2 Auction-based Approaches to Negotiation in Highly Rugged Utility Spaces

Ito et al. [4] presented a bidding-based protocol to deal with nonlinear utility spaces generated using weighted constraints. The protocol consists on the following four steps:

1. *Sampling:* Each agent takes a fixed number of random samples from the contract space, using a uniform distribution.

2. *Adjusting:* Each agent applies simulated annealing to each sample to try to find a local optimum in its neighborhood. This results in a set of high-utility contracts.

3. *Bidding:* Each agent generates a bid for each high-utility, adjusted contract. The bids are generated as the intersection of all constraints which are satisfied by the contract. Each agent sends its bids to the mediator, along with the utility associated to each bid.

4. *Deal identification:* The mediator employs breadth-first search with branch cutting to find overlaps between the bids of the different agents. The regions of the contract space corresponding to the intersections of at least one bid of each agent are tagged as potential solutions. The final solution is the one that maximizes

joint utility, defined as the sum of the utilities for the different agents.

In [8], we proposed an alternative perspective for the bidding process, looking at the constraint-based agent utility space as a weighted undirected graph. Consider again the simple utility space example shown in Figure 1. Think about each constraint as a node in the graph, with an associated weight which is the utility value associated to the constraint. Now we will connect all nodes whose corresponding constraints are *incompatibles*, that is, they have no intersection. The resulting graph is shown in Figure 2.

To find the highest utility bid in such a graph can be seen as finding the set of unconnected nodes which maximizes the sum of the nodes' weights. Since only incompatible nodes are connected, the corresponding constraints will have non-null intersection. In the example, this would be achieved by taking the set $\{C1, C2\}$. The problem of finding a maximum weight set of unconnected nodes is a well-known problem called maximum weight independent set (MWIS). Though MWIS problems are NP-hard, in [1], a message passing algorithm is used to estimate MWIS, which greatly reduces the complexity of the search.

Since the algorithm is deterministic, only one bid can be generated for a given set of constraints. To solve this, in [8], the algorithm is applied to a subset of constraints $C' = \{c'_k | k = 1, ..., n_c; n_c < l; c'_k \in C\}$. The constraints c'_k are randomly chosen from the constraint set C. In this way, a different constraint subset C' is passed to the algorithm at each run, which will result in different, non-deterministic bids.

Both approaches are evaluated in nonlinear scenarios for different number of agents and issues, and they achieve great results in terms of optimality (measured as the ratio between the solutions found using the protocol and the optimal solution computed using complete information) and failure rate (measured as the ratio between unsuccessful negotiations and total negotiations).

2.3 Constraint/Bid Quality Factor for Highly Rugged Utility Spaces

The use of weighted constraints generates a "bumpy" utility space, with many peaks and valleys. However, the degree of "bumpiness" may vary from one scenario to another. More formally, the complexity of the utility spaces of the agents may be measured using a correlation length, which has been widely used to asses fitness landscape complexity in evolutionary computation [13]. Correlation length is defined as the minimum distance between samples in the utility space which makes the correlation between those samples drop below a given threshold. Intuitively, the main difference between highly correlated and highly uncorrelated (i.e. highly rugged) utility spaces is the width of the high-utility regions. Highly-rugged scenarios will yield narrower peaks. Since the mechanisms outlined above lead agents to choose those peaks (or high-utility regions) as bids, the result is that narrower bids will be sent to the mediator. The width of the bids (or more generally, the *volume* of the bids, computed as the cardinality of the set of contracts which match the bid), will directly impact the probability that the bid overlaps a bid of another agent, and thus its *viability*, that is, the probability of the bid resulting in a deal. In [9], we introduced the hypothesis that an agent with no knowledge of the other agents' preferences should try to adequately balance the utility of their bids (to maximize its own profit) and the volume of those bids (to maximize the probability of a successful negotiation). To allow this, we defined the *quality factor* of a constraint or a bid as $Q_c = u_c^{\alpha} \cdot v_c^{1-\alpha}$, where u_c and v_c are, respectively, the utility and volume of the bid or constraint c, and $\alpha \in [0, 1]$ is a parameter which models the risk attitude of the agent. A risk averse agent ($\alpha < 0.5$) will tend to qualify as better bids those that are wider, and thus are more likely to result in a deal. A risk willing or selfish agent ($\alpha > 0.5$) will, in contrast, give more importance to bid utility. Finally, we proposed a set of mechanisms to integrate this quality factor into the bidding and deal identification steps of the negotiation process described in their previous works, and validate their hypothesis through a set of experiments that show how the proposed approach improves the negotiation process both in terms of effectiveness and performance.

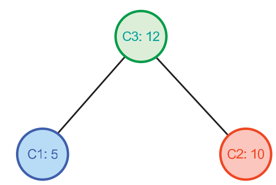

Figure 2: Weighted undirected graph resulting from the utility space in Figure 1

Though the approach proposed in [9] yields satisfactory results in highly-rugged utility spaces, there are some issues which are not addressed in the work. Even when the quality factor is designed to model the risk attitudes of the agents through its α parameter (and thus α allows to model agent strategies), the experimental evaluation is performed only for $\alpha = 0.5$. This assumes that all negotiating agents have the same attitude towards risk, and also that this risk attitude is neutral (i.e. agents give the same weight to utility and deal probability). In a real, competitive scenario, these assumptions do not necessarily hold, and therefore a strategy analysis is needed to evaluate the mechanisms in situations where agents with different risk attitudes interact.

3. STRATEGY ANALYSIS

One of the main challenges in mechanism design for automated negotiation is strategic stability. Here stability is seen as the impossibility (or at least difficulty) to manipulate the mechanisms by means of strategies. This means that the mechanisms must motivate the agents to behave in an adequate way, since if a rational agent may benefit from taking a given strategy instead of the one expected by the mechanisms, it will do so. This problem is closely related to the notion of *equilibrium* in game theory [14, 6]. For heuristic approaches such as those described in [9], game theory analyses cannot be directly applied, but some of the concepts can still be useful.

For instance, we can talk about a *dominant strategy* if there is a strategy which is always the best choice for an agent, whatever the other agents do. In most cases, however, an agent's best strategy depends on the strategies used by its opponents, and stability is achieved by means of strategy profiles (sets of strategies for every agent). A given strategy profile $F = \{f_1, \ldots, f_N\}$ is said to be an *equilibrium* for a given setting if every agent i has no better strategy than f_i, provided that the other agents j play their corresponding strategies f_j. That means that, if all agents use their corresponding equilibrium strategies, there is no incentive for any agent to deviate from that set of strategies [10]. Of course, finding these equilibrium conditions for a given mechanism may be a very complex task. Furthermore, equilibrium conditions may be non-unique, which generates the additional problem of determining which of the equilibrium strategy profiles to use for a given negotiation. Finally, the stability of a negotiation mechanism does not guarantee social welfare maximizing solutions. In many situations, individual rationality drives agents towards strategies which yield low individual and social welfares. Those situations, which should be avoided in mechanisms design, are usually called instances of the well-known *prisoner's dilemma* [12] or, more generally, situations of high *price of anarchy (PoA)* [11].

As we stated above, the equilibrium concepts outlined here are related to game theory, and thus are very difficult to determine for heuristic negotiation mechanisms. However, probabilistic analysis and empirical evaluations can be performed in an analogous manner for these mechanisms. The rest of this section is dedicated to assess the strategic behavior of the approach proposed in [9], determining the existence of dominant individual strategies and social optimal strategies, and verifying if the auction-based negotiation mechanisms are prone to the prisoner's dilemma.

3.1 Probabilistic Analysis

Intuitively, it can be seen that the quality factor defined in [9] allows an agent to balance bid utility (to maximize its own benefit) and bid volume (to maximize deal probability). More formally, we may find mathematic expressions for the deal probability and the expected utility in a negotiation using the auction-based protocol. The deduction of these expressions is beyond the scope of this paper, and can be found in [7]. For this work, the final expressions will suffice. In particular, deal probability for a single run of the auction-based negotiation protocol is given by

$$P_{deal} = \sum_{j=1}^{\prod n_{bp}^k} (-1)^{j+1} \binom{\prod n_{bp}^k}{j} \left(\frac{1}{|D|^{n(n^a-1)}} \right)^j, \quad (1)$$

where n^a is the number of negotiating agents, n is the number of issues, $|D|$ is the domain size for the issues (assuming all issues have the same domain size), and n_{bp}^k is the *number of bidden contracts* for agent k, that is, an indication of the portion of the solution space which is covered by agent k bids. This is given by $n_{bp}^k = \sum_{l=1}^{n_b^k} v_l^k$, where n_b^k is the number of bids issued by agent k and v_l^k is the volume of each l-th bid.

In a similar way, we can see that the *expected utility* for agent k is given by

$$E[u^k] = \left[\sum_{l=1}^{n_b^k} u_l^k \cdot v_l^k \right] \left[\sum_{j=1}^{\prod n_{bp}^k} \binom{\prod n_{bp}^k}{j} \frac{(-1)^{j+1}}{|D|^{n(n^a-1)j}} \right], \quad (2)$$

where u_l^k is the utility for the l-th bid of agent k. According to this expression, to maximize expected utility, an agent should reveal as much information as possible. If information disclosure is limited, an agent should try to maximize $\sum_{l=1}^{n_b^k} u_l^k \cdot v_l^k$, balancing in this way bid utility and bid volume. This is coherent to the choice of $\alpha = 0.5$ in [9]. Of course, this strategy does not model the attitude of a risk willing agent, who would prefer to risk the success of the negotiation to have the chance of a higher utility gain. To model this, we can use an *expected deal utility*, that is, the expected utility for an agent provided that a deal has been reached. This expected deal utility is given by:

$$E[u^k \,|deal] = \frac{\sum_{l=1}^{n_b^k} u_l^k \cdot v_l^k}{n_{bp}^k} \quad (3)$$

According to this, a risk willing agent would give preference to bid utility against bid volume, trying to reduce n_{bp}^k to maximize expected deal utility, but reducing also deal probability.

These expressions are coherent with the intuitive notion of agent risk attitude introduced in the quality factor in [9]. We can also use them to infer some of the strategic properties of the protocol. Since deal probability increases with deal volume, low values of α are expected to increase deal probability too. As we have seen, when there is total uncertainty about the utility spaces of the agents, the expected utility is maximized for $\alpha = 0.5$. If the utility spaces of the agents are specially complex (highly uncorrelated), it is reasonable to think that the deal probability will be lower, and thus agents should use lower values of α (that is, they should take less risks) in order to keep expected utility at an acceptable value. Similarly, if the agent's utility spaces are highly correlated, agents could use higher α values (that is, be more utility oriented), trying to maximize the expected deal utility, since deal probability will be higher. Furthermore, since lower α values increase deal probability, a single agent could benefit from a risk-willing strategy if the other agents are risk averse (their lower α values would compensate the decrement in deal probability). However, should all agents decide to use risk willing strategies, deal probability would reduce drastically, leading to low expected individual and social welfares. As stated above, this would be a situation of high price of anarchy, analogous to the prisoner's dilemma.

3.2 Experimental Analysis

In this section the strategic properties of the protocol inferred from the statistical analysis are empirically verified. First of all, individual equilibrium is studied, trying to determine the existence of a *dominant strategy*, which is the best strategy for an agent under any circumstance, or whether there is an optimal strategy for an agent depending on the strategies of the other agents.

To evaluate this, we have performed a set of experiments comparing the utility obtained by a *individualist agent*, using a strategy determined by α_i, with the utility obtained by the other agents, which use a strategy determined by α_s.

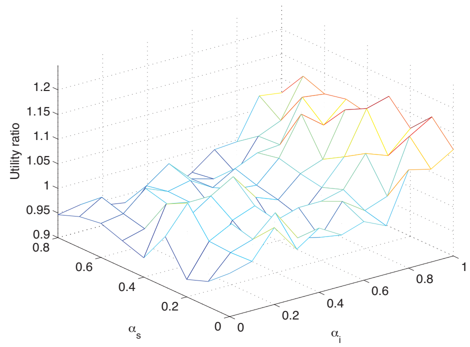

Figure 3: Individual equilibrium analysis

Experiments have been performed varying α_i y α_s within the interval $[0, 1]$ in 0.1 steps.

Figure 3 shows the results of the experiments for 6 agents and 6 issues. We have represented the ratio between the utility obtained by the individualist agent and the utility obtained by the rest of the agents for different values of α_i y α_s. For $\alpha_s > 0.8$ negotiations failed, and thus no values are shown in the figures. Generally, the individualist agent obtains a higher utility than the rest of the agents when using higher α_i values. In particular, for any α_s, the maximum utility value for the individualist agent is obtained for $\alpha_i = 1$, so we can conclude that this is the individual *dominant strategy* for the agents.

Once individual strategies have been analyzed, we have studied social strategies, trying to determine the existence of a set of strategies for the different agents which maximizes social welfare. Since both the negotiation model and the measure we have taken for social welfare (Nash product) are symmetric, we expect this strategy set to be symmetric as well. Taking this into account, we have performed a set of experiments using for all agents the same *social strategy*, determined by α_s. Experiments have been conducted varing α_s within the interval $[0, 1]$ in 0.1 steps. Furthermore, to study the variation of the results with the complexity of the utility spaces, the experiments have been repeated for utility spaces of different complexity. Utility space complexity have been measured using a correlation length ψ. As introduced in Section 3.1, correlation length is defined as the minimum distance between samples in the utility space which makes the correlation between samples drop below a given threshold. For the purpose of this work, we have chosen a threshold of 0.7.

Experiment results for 6 agents and 6 issues are presented in Table 1. The table shows the median social welfare optimality for the negotiation as the value of α_s varies, for different values of ψ. Optimality is defined as the ratio between the social welfare obtained with the protocol and the social welfare obtained using an optimizer with complete information. We can see that the α values which maximize social optimality are around 0.6 and 0.8. This is higher than the theoretical optimum ($\alpha = 0.5$), which is reasonable if we think that calculations were made assuming total uncertainty about the utility space (that is, $\psi = 0$).

Once an optimal social strategy has been identified, a desirable property would be that this strategy were an equilibrium for the system, that is, that there was no incentive for any agent to deviate from this strategy. Unfortunately, as we seen above, there is a dominant individual strategy, given by $\alpha_i = 1$. Therefore, an individually rational agent may decide to take this strategy to maximize its own benefit. All agents have the same incentive, so equilibrium is reached when all agents choose $\alpha_i = 1$. As we can see in Table 1, this makes negotiation fails in medium and highly complex scenarios. This confirms that the protocol is prone to the prisoner's dilemma.

4. AVOIDING THE PRISONER'S DILEMMA IN THE AUCTION-BASED NEGOTIATION PROTOCOL

In this section, stability problems of the auction-based negotiation protocol are addressed. A set of different mechanisms intended to avoid high price of anarchy situations in the negotiation process are proposed, and their effectiveness is empirically evaluated.

4.1 Enforcing Socially-oriented Strategies at the Mediator

To improve the strategic stability of the negotiation, the mechanisms should be modified to incentivize the adoption of socially optimal strategies. The logical step in the protocol to make any modification is the deal identification at the mediator. Since negotiating agents are supposed to be individually rational, it is the mediator the only one who can be assumed to pursue social welfare. In the basic protocol proposed in [9], the mediator chooses as the final solution the one maximizing social welfare, computed as the Nash product of the individual agent utilities. Since the Nash product is symmetric, those agents whose bids have higher average utility would, on average, obtain higher utilities in the final deal, which incentivizes the use of the dominant strategy. To mitigate this effect, a reasonable measure could be to reward in the selection of the final solution to those agents which have made wider bids. To achieve this, we propose a modification of the Nash product which we have called *weighted product by average volume*:

$$sw_{\bar{V}}(s, U) = \prod_{i=1}^{n_a} \left(u^i\left(s\right)\right)^{\frac{\bar{v}^i}{\max_{1 \le j \le n_a} \bar{v}^j}}, \qquad (4)$$

where $u^i(s)$ is the utility of the solution s for agent i, and $\bar{v}^i$ is the average volume of the bids issued by agent i.

In this way, the utility for those agents who have issued widest bids (which, on average, will be the ones using more socially oriented strategies) will be given more weight in the selection of the final solution than those of the more selfish agents. An interesting effect of this metric is that a rational agent could issue some high volume, low utility bids to try to compensate for its high-utility, low volume bids. To counter this effect, a *product weighted by average quality factor* is proposed:

$$sw_{\bar{Q}}(s, U) = \prod_{i=1}^{n_a} \left(u^i\left(s\right)\right)^{\frac{\bar{Q}^i}{\max_{1 \le j \le n_a} \bar{Q}^j}}, \qquad (5)$$

where $\bar{Q}^i$ is the average quality factor of the bids issued by agent i.

Table 1: Social strategy analysis

		α_s										
		0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
$\psi_{0.7}$	2.8	0.3335	0.3788	0.3836	0.3765	0.4336	0.4801	0.5521	0.4855	0	0	0
	3.1	0.4600	0.5282	0.4951	0.5041	0.5544	0.5553	0.5960	0.6822	0	0	0
	4.0	0.7954	0.7849	0.7977	0.8137	0.8211	0.8380	0.8283	0.8270	0.8139	0	0
	4.3	0.9672	0.9634	0.9759	0.9608	0.9728	0.9690	0.9710	0.9707	0.9774	0	0
	4.6	1.0000	1.0000	0.9748	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
	5.9	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

Finally, bid selection for deal identification at the mediator is performed using the quality factor of the bids *as declared by the agent issuing the bids*. This makes the assesment of the bids made by the mediator dependent on the risk attitudes of the agents, thus favoring those agents with more selfish strategies. Taking this into account, we propose that the mediator uses its own α_m parameter for Q calculation. In this way, we expect to decouple deal identification from the negotiating agent strategies, improving the stability of the protocol. Possible choices for α_m are the optimal social strategy for a given correlation length, or $\alpha_m = 0.5$, which is the theoretical optimal value if there is total uncertainty about the agents' utility spaces. However, there is a problem with using such α_m values. Any $\alpha_m \geq 0.5$ would give at least the same weight to bid utility than to bid volume. Because of this, it would not be possible for the mediator to discriminate whether a given bid has a high quality factor due to its high volume (thus being probably a bid issued by a socially oriented agent) or due to its high utility (and thus probably generated by a selfish agent). It seems reasonable to use $\alpha_m < 0.5$, giving more weight to higher volume bids, and thus enforcing social behavior among agents. The limit would be to use $\alpha_m = 0$, which would make the mediator to select bids according only to their volume, regardless of their utility. Our hypothesis is that this would totally decouple the deal identification mechanism from the strategic behavior of the negotiating agents, thus improving protocol stability.

4.2 Stability Analysis

Stability analysis is oriented to determine the possibility of an agent manipulating the negotiation to its own benefit. In the model we are dealing with, this manipulation may occur when an agent deviates from the social strategy taking a more selfish approach. To evaluate this empirically, we have performed experiments comparing the utility obtained by an *individualist agent*, using its dominant strategy $\alpha_i = 1$, against the utility obtained by the rest of agents, which will be using the corresponding optimal social strategy α_s. Experiments have been made for utility spaces with different correlation lengths. Furthermore, since the model is designed for multiagent negotiations, experiments have been performed for different number of individualist agents, thus studying the effect of possible coalitions or coincidences of selfish agents.

Table 2 presents the results of the experiments for 6 agents and 6 issues, showing the ratio between utilities for individualist agents and social agents for different correlation lengths and different number of individualist agents. The table shows the medians and the 95% confidence intervals for 100 runs of each experiment. We can see that there is only a significative gain for the individualist agents in medium complexity scenarios. For highly complex scenarios (low correlation length), individualism make negotiations fail, and thus there is no incentive for any agent to deviate from the social strategy. For more correlated scenarios ($\psi_{0.7} = 4$), we can see that an individualist agent may obtain a benefit of about 200%, though coalitions are not likely, since an increase in the number of individualist agents make negotiations fail. For $\psi_{0.7} = 4.3$, coalitions seem viable, as they increase the benefit for the individualist agents. Finally, for the less complex scenarios ($\psi_{0.7} \geq 5.9$), a selfish attitude does not provide a significant difference in utility, since all agents get high utilities with the social strategy. From these results we can conclude that the model is stable in low complexity and high complexity scenarios, and that the scenarios of medium complexity make stability problems arise, and require the application of additional mechanisms.

In the previous section, a set of alternative mechanisms for deal identification at the mediator were proposed. Those mechanisms were intended to incentivize agents to social behavior, and thus solve the stability problems of the model. To evaluate the effect of the proposed mechanisms on the stability of the protocol, we have repeated the experiments for the different approaches discussed in Section 4.1:

Nash Reference approach, using Nash product.

Average_V Product weighted by average bid volume (Equation 4).

Average_$Q_{0.5}$ Product weighted by average quality factor (Equation 5), with $\alpha_m = 0.5$, corresponding to the theoretical optimal social strategy.

Average_Q_0 Product weighted by average quality factor (Equation 5), with $\alpha_m = 0$, corresponding to a deal identification strategy totally decoupled from agent utility (the mediator only considers bid volume).

Figure 4 presents the results of the experiments for 6 agents and 6 issues for the most critical scenarios in terms of stability found in the previous experiment ($\psi_{0.7} = 4$ and $\psi_{0.7} = 4.3$). Each graphic presents a box-plot for the final outcomes of 100 runs of the experiment. The horizontal axis represents the approach under evaluation, while in the vertical axis we have represented the optimality rate as notched box and whisker plots. The boxes have lines for the median and the 25th and 75th percentiles of the gain for individualist agents in each negotiation (computed as the ratio between the utilities obtained by individualist and social agents), and the whiskers show adjacent values in the data. Outliers are displayed with a plus (+) sign. Notches display the variability of the median between samples. We can see that the mechanism based on average volume provides no improvement is stability, since for both cases median utility results are higher for individualist agents. The mechanism based on average quality factor, however, adequately improves the stability of the protocol, and this improvement is greater for $\alpha_m = 0$. From these results we can conclude that decoupling deal identification from the attitudes of the

Table 2: Stability analysis for 6 agents and 6 issues

		Number of individualist agents					
		1		2		3	
		median	conf. interval	median	conf. interval	median	conf. interval
$\psi_{0.7}$	2.8	–	–	–	–	–	–
	3.1	–	–	–	–	–	
	4.0	2.0086	[1.8574,2.1598]	–	–	–	–
	4.3	1.1066	[1.0610, 1.1522]	1.1986	[1.1431,1.2541]	–	–
	4.6	0.9795	[0.9567, 1.0024]	1.0081	[0.9870, 1.0292]	0.9785	[0.9567,1.0003]
	5.9	1.0336	[1.0081, 1.0591]	1.0243	[1.0043, 1.0443]	0.9811	[0.9598,1.0024]

negotiating agents by making the mediator calculate its own quality factor improves strategic stability of the negotiation process.

Finally, we need to consider the effect of these approaches over the Nash optimality of the negotiation. Since these techniques give preference to socially oriented offers against higher utility offers, this may make final deals to be further from the theoretical optimum. To evaluate this, Table 3 show median optimality rates for the different approaches, measured as the ratio between the Nash utility achieved by the protocol and the Nash utility obtained using an optimizer with complete information. As a reference, the results obtained by the approach in [9] have been included. Results show that the approaches which improve stability suffer a slight decrement in optimality for the most correlated scenarios, and an increment in optimality for the most uncorrelated ones (due to the decrement in failure rate). However, this decrement is not statistically significative. We can conclude that it is possible to stabilize the model to a great extent by having the mediator compute its own quality factor Q, and that this improved stability has no significative impact over Nash optimality.

5. DISCUSSION AND RELATED WORK

There is a number of recent research lines addressing complex negotiations. Most of them provide a way to overcome the complexity of intractable utility spaces, be it by approximating these complex spaces by means of simpler utility functions [3], or by developing heuristic mechanisms which perform a more efficient search for deals in the solution space [4]. However, very few works address agents' strategic behavior in their proposed models. In [9], we find the first reference in complex negotiations to the possibility of agents having a wide range of strategies with the notion of risk attitude they introduce in their model. Though the model supports agents' strategic behavior, they do not analyze the negotiation dynamics when agents with different strategies interact, neither prove the strategic stability of their model. In this paper, we have performed both a theoretical analysis and an experimental evaluation of the model. This hybrid approach is motivated by the fact that the decision mechanisms for the negotiating agents and the mediator are based on heuristics, and thus a systematic analysis as the performed in [2] was not feasible.

The strategy analysis has allowed us to identify some serious stability concerns, like the fact that the auction-based protocol is prone to high price of anarchy situations, which is an undesirable property for any negotiation approach. To overcome this problem, we have proposed a set of measures intended to incentivize social behavior among negotiating agents. The proposed mechanisms are based on biasing deal identification at the mediator to give preference to the most socially oriented bids. This is somewhat similar to the approach taken in [5] for bilateral negotiations with binary is-

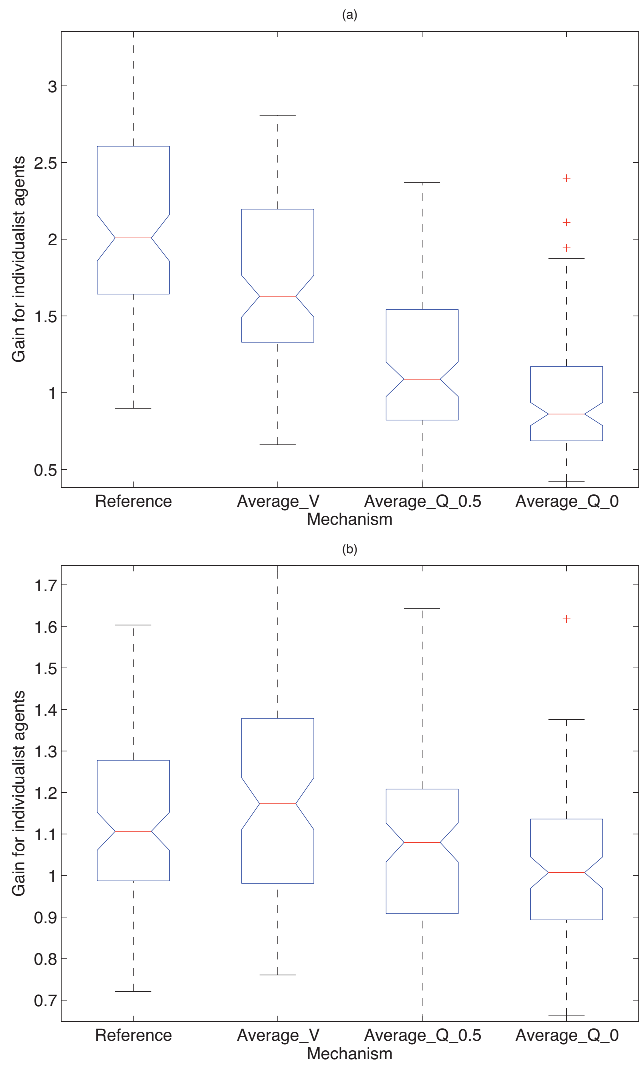

Figure 4: Effect of the different mechanisms on the stability of the protocol for the most critical scenarios: a) $\psi_{0.7} = 4$, b) $\psi_{0.7} = 4.3$

Table 3: Effect of the different mechanisms over Nash optimality

		Mechanism							
		Reference		Average_V		Average_Q_0.5		Average_Q_0	
		median	conf. interval	median	conf. interval	median	conf. interval	median	conf. interval
$\psi_{0.7}$	2.8	0	[0, 0]	0.5028	[0.4696, 0.5359]	0.5358	[0.5043, 0.5674]	0.5196	[0.4888, 0.5503]
	3.1	0	[0, 0]	0.5828	[0.5546, 0.6110]	0.6063	[0.5783, 0.6343]	0.5962	[0.5727, 0.6196]
	4.0	0.7733	[0.7319, 0.8148]	0.8379	[0.8092, 0.8667]	0.7731	[0.7509, 0.7954]	0.8142	[0.7900, 0.8384]
	4.3	0.9746	[0.9521, 0.9808]	0.9639	[0.9580, 0.9698]	0.9617	[0.9539, 0.9694]	0.9734	[0.9655, 0.9812]
	4.6	1.0000	[1.0000, 1.0000]	1.0000	[0.9951, 1.0000]	1.0000	[0.9953, 1.0000]	1.0000	[0.9936, 1.0000]
	5.9	1.0000	[1.0000, 1.0000]	1.0000	[1.0000, 1.0000]	1.0000	[1.0000, 1.0000]	1.0000	[1.0000, 1.0000]

sue dependencies, though our approach allows the negotiating agents to retain control of their strategic profile, instead of delegating it to the mediator.

6. CONCLUSIONS AND FUTURE WORK

The prisoner's dilemma, or more generally, high price of anarchy situations, which imply that individual rationality drives the agents towards strategies which yield low individual and social welfares, are conditions which should be avoided at all costs when designing negotiation mechanisms. This is specially important when dealing with complex negotiations involving highly rugged utility spaces, since in these cases "low individual and social welfare" often means that the negotiations fail. Therefore, an strategic analysis is paramount for any model intended to work for highly rugged utility spaces, in order to determine the strategic properties of the model and to allow to establish additional mechanisms for stability if needed.

In this paper we have performed a strategy analysis for the auction based negotiation protocol for highly rugged utility spaces proposed in [9]. This strategy analysis has started studying the equilibrium conditions, which has revealed the existence of an individual dominant strategy, which is different from the socially optimal strategy. A more in-depth stability analysis has shown that, for highly correlated or lowly correlated scenarios, there is no incentive for negotiating agents to deviate from the socially optimal strategy. However, for medium complexity scenarios a selfish agent may benefit from using its dominant strategy, which raises stability concerns, leading the model to a situation analogous to the well-known prisoner's dilemma. To solve this, we have proposed a set of mechanisms intended to incentivize social behavior among negotiating agents. These mechanisms are based on biasing deal identification at the mediator towards those bids which are more socially oriented, thus decoupling the search for social welfare from the individual agents' goals. Experiments show that the proposed mechanisms successfully stabilize the protocol. However, there is still plenty of research to be done in this area. We are interested on adaptive measures, allowing the mediator to deduce agent strategies during the negotiation process, and thus to apply the different mechanisms as needed. In addition, the effect of the correlation between the utility functions of *different* agents (as opposed to the correlation distance within each agent's utility function) should be analyzed. Finally, we are working on the generalization of these approaches for other negotiation protocols and utility function types.

7. ACKNOWLEDGMENTS

This work has been partially supported by the Spanish Ministry of Education and Science grant TIN2008-06739-C04-04.

8. REFERENCES

[1] M. Bayati, D. Shah, and M. Sharma. Max-product for maximum weight matching: Convergence, correctness, and lp duality. *IEEE Transactions on Information Theory*, 54(3):1241–1251, 2008.

[2] S. S. Fatima. Bidding strategies for multi-object auctions. *Negotiation, Auctions, and Market Engineering*, pages 200–212, 2008.

[3] K. Hindriks, C. M. Jonker, and D. Tykhonov. Eliminating interdependencies between issues for multi-issue negotiation. In *Cooperative Information Agents X*, volume 4149 of *Lecture Notes in Computer Science*, pages 301–316, 2006.

[4] T. Ito, M. Klein, and H. Hattori. Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI07)*, pages 1347–1352, 2007.

[5] M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam. Protocols for negotiating complex contracts. *IEEE Intelligent Systems*, 18(6):32–38, 2003.

[6] S. Kraus. Automated negotiation and decision making in multiagent environments. In *Mutli-agents systems and applications*, pages 150–172. Springer-Verlag New York, Inc., New York, NY, USA, 2001.

[7] I. Marsa-Maestre. *Contribución a la negociación automática en espacios de utilidad complejos*. PhD thesis, Universidad de Alcalá, 2009.

[8] I. Marsa-Maestre, M. A. Lopez-Carmona, J. R. Velasco, and E. de la Hoz. Effective bidding and deal identification for negotiations in highly nonlinear scenarios. In *Proc. of the 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, pages 1057–1064, 2009.

[9] I. Marsa-Maestre, M. A. Lopez-Carmona, J. R. Velasco, T. Ito, K. Fujita, and M. Klein. Balancing utility and deal probability for negotiations in highly nonlinear utility spaces. In *Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI-09)*, pages 214–219, 2009.

[10] J. F. Nash. Two-person cooperative games. *Econometrica*, 21(1):128–140, 1953.

[11] C. Papadimitriou. Algorithms, games, and the internet. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 749–753, New York, NY, USA, 2001. ACM.

[12] W. Poundstone. *Prisoner's Dilemma*. Anchor, New York, USA, 1993.

[13] E. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63(5):325–336, 1990.

[14] G. Weiss. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, 1999.