

Fig. 8 Stability analysis of the protocol using SA-Q for scenarios with different correlation lengths. **a** $\psi = 2.8$, **b** $\psi = 4$, **c** $\psi = 4.3$, **d** $\psi = 5.9$

scenario induced by individually rational combinations of strategies, yielding zero utility for all agents, which imply an infinite expected price of anarchy (EPoA). This is an undesirable property of the model, and requires 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 Sect. 5.1:

- *Nash*: Reference approach, using Nash product.
- *Average_V*: Product weighted by average bid volume (Eq. 4).
- *Average_Q_{0.5}*: Product weighted by average quality factor (Eq. 5), with $\alpha_m = 0.5$, corresponding to the theoretical socially optimal strategy. This α_m is also used for deal identification at the mediator, as described in Sect. 5.1.
- *Average_Q₀*: Product weighted by average quality factor, with $\alpha_m = 0$, corresponding to a deal identification strategy totally decoupled from agent utility (the mediator only considers bid volume). This α_m is also used for deal identification at the mediator.

Figures 10 and 11 present the results of the experiments for SA-Q and MWIS-Q, respectively. The figures show the results for 6 agents and 6 issues with utility spaces of correlation

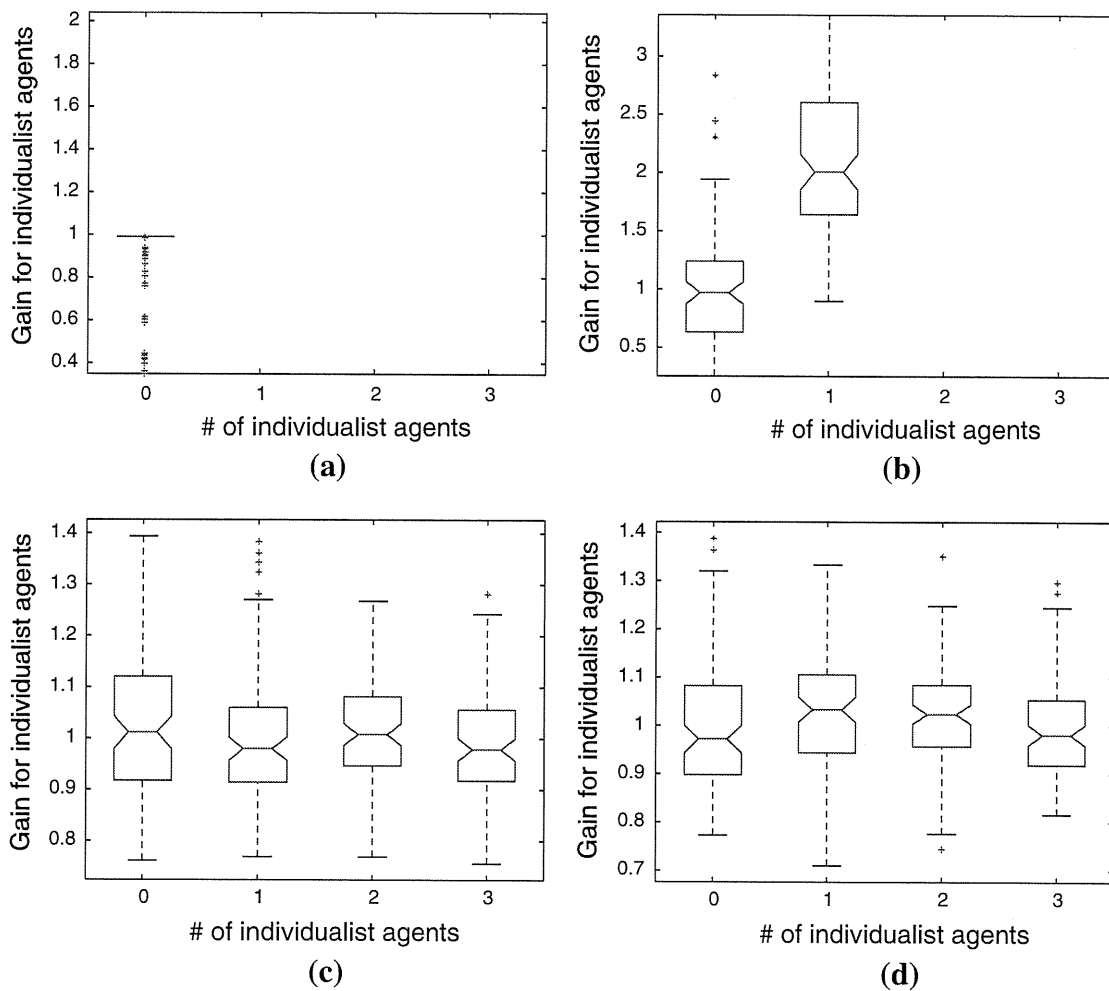


Fig. 9 Stability analysis of the protocol using MWIS-Q for scenarios with different correlation lengths. **a** $\psi = 2.8$, **b** $\psi = 4$, **c** $\psi = 4.3$, **d** $\psi = 5.9$

Table 3 Stability analysis for SA-Q with six agents and six issues: gain for individualist agents against social agents

ψ	Number of individualist agents					
	1		2		3	
	median	conf. interval	median	conf. interval	median	conf. interval
2.8	–	–	–	–	–	–
3.1	–	–	–	–	–	–
4.0	1.5440	[1.4170, 1.6710]	–	–	–	–
4.3	1.0861	[1.0436, 1.1286]	1.1203	[1.0902, 1.1503]	1.1648	[1.1207, 1.2090]
4.6	0.9993	[0.9663, 1.0322]	1.0208	[0.9981, 1.0435]	1.001	[0.9796, 1.0224]
5.9	0.9693	[0.9438, 0.9949]	0.9976	[0.9775, 1.0177]	0.9907	[0.9715, 1.0100]

lengths $\psi = 4$ and $\psi = 4.3$, which were identified in the previous experiment as the most critical scenarios regarding stability. 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 gain for individualist agents in each

Table 4 Stability analysis for MWIS-Q with six agents and six issues: gain for individualist agents against social agents

ψ	Number of individualist agents					
	1		2		3	
	median	conf. interval	median	conf. interval	median	conf. interval
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]

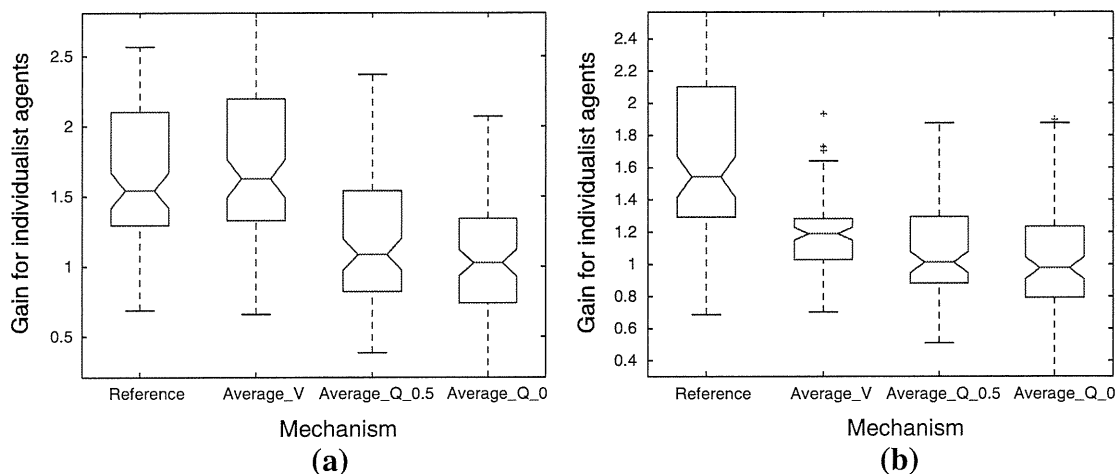


Fig. 10 Effect of the different mechanisms on the stability of the protocol for SA-Q in the most critical scenarios. **a** $\psi = 4.0$, **b** $\psi = 4.3$

negotiation. We can see that the mechanism based on average volume provides not enough improvement in stability, since for all cases median utility results are higher for selfish agents, thus maintaining the incentive for agents to deviate from the socially optimal strategy. The mechanism based on average quality factor, however, significantly mitigates the gain for selfish agents, removing the incentive to choose the previously individually optimal strategy ($\alpha = 1$). Due to the effect of this mechanism, the situation where all agents take selfish strategies is no longer induced by individually rationality, thus avoiding the infinite Expected Price of Anarchy values. This 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 negotiating agents by making the mediator calculate its own quality factor improves the strategic stability of the negotiation process, significantly decreasing Expected Price of Anarchy.

Since the 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, as discussed in Sect. 2.1.3, we can consider the Price of Stability (PoS) imposed by the proposed mechanisms. As it occurred with PoA, we cannot use Price of Stability definition directly,

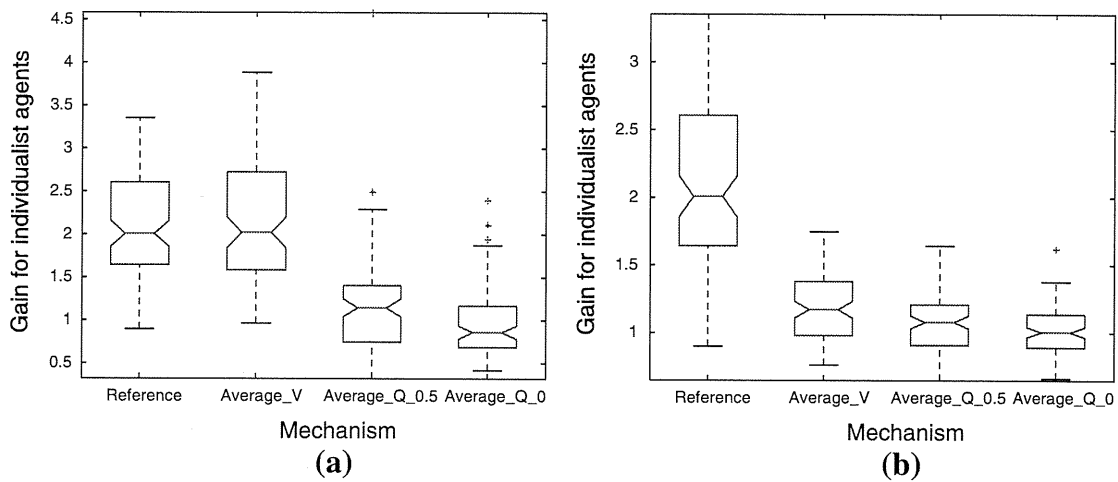


Fig. 11 Effect of the different mechanisms on the stability of the protocol for MWIS-Q in the most critical scenarios. **a** $\psi = 4.0$, **b** $\psi = 4.3$

Table 5 Effect of the different mechanisms over social optimality rate (and thus, over expected price of stability) for SA-Q

ψ	Mechanism							
	Reference		Average_V		Average_Q_0.5		Average_Q_0	
	Med.	Conf. interval	Med.	Conf. interval	Med.	Conf. interval	Med.	Conf. interval
2.8	0.326	[0.305, 0.347]	0.633	[0.608, 0.658]	0.614	[0.587, 0.641]	0.632	[0.609, 0.655]
3.1	0.530	[0.511, 0.549]	0.588	[0.562, 0.614]	0.572	[0.546, 0.598]	0.571	[0.545, 0.597]
4.0	0.769	[0.740, 0.798]	0.620	[0.583, 0.657]	0.600	[0.566, 0.634]	0.625	[0.594, 0.656]
4.3	0.960	[0.951, 0.969]	0.734	[0.696, 0.771]	0.756	[0.715, 0.797]	0.727	[0.688, 0.767]
4.6	1.000	[1.000, 1.000]	0.836	[0.811, 0.860]	0.856	[0.831, 0.881]	0.847	[0.826, 0.869]
5.9	1.000	[1.000, 1.000]	0.939	[0.922, 0.956]	0.953	[0.938, 0.967]	0.967	[0.953, 0.981]

since it relies on Nash equilibrium conditions. We can, however, define *Expected Price of Stability (EPoS)* in an analogous way as we defined EPoSA in the previous section:

Definition 15 *Expected Price of Stability (EPoS)* The *Expected Price of Stability* in a non-deterministic game is the ratio between the maximum expected social welfare achievable by means of a feasible agent strategy combination and the maximum expected social welfare achievable by means of an *individually rational* agent strategy combination.

$$EPoS = \frac{\max_{s \in S} E[sw(s)]}{\max_{s \in S_{i.r.}} E[sw(s)]}$$

where S is the set of all feasible strategy combinations of the game, $S_{i.r.} \subseteq S$ is the set of all strategic combinations which are individually-rational for the negotiating agents, and $E[sw(s)]$ is the expected social welfare for a given strategy combination s .

Tables 5 and 6 present the median social optimality rates for SA-Q and MWIS-Q, respectively, using the different mechanisms proposed, when all negotiating agents choose the socially optimal strategy. The statistic on this ratio is analogous the inverse of the Expected Price of Stability defined above. As a reference, the results obtained when no asymmetrical

Table 6 Effect of the different mechanisms over Social Optimality Rate (and thus, over Expected Price of Stability) for MWIS-Q

ψ	Mechanism							
	Reference		Average_V		Average_Q_0.5		Average_Q_0	
	Med.	Conf. interval	Med.	Conf. interval	Med.	Conf. interval	Med.	Conf. interval
2.8	0.553	[0.521, 0.585]	0.503	[0.470, 0.536]	0.536	[0.504, 0.567]	0.520	[0.489, 0.550]
3.1	0.681	[0.652, 0.710]	0.583	[0.555, 0.611]	0.606	[0.578, 0.634]	0.596	[0.573, 0.620]
4.0	0.949	[0.925, 0.973]	0.838	[0.809, 0.867]	0.773	[0.751, 0.795]	0.814	[0.790, 0.838]
4.3	0.975	[0.952, 0.981]	0.964	[0.958, 0.970]	0.962	[0.954, 0.969]	0.973	[0.966, 0.981]
4.6	1.000	[1.000, 1.000]	1.000	[0.995, 1.000]	1.000	[0.995, 1.000]	1.000	[0.994, 1.000]
5.9	1.000	[1.000, 1.000]	1.000	[1.000, 1.000]	1.000	[1.000, 1.000]	1.000	[1.000, 1.000]

social welfare metric is used have been included. Results show that, for SA-Q, the approaches which improve stability suffer a significant 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). For MWIS-Q a similar trend is observed, though the optimality loss is lower. We can conclude that, though it is possible to stabilize the model to a great extent by having the mediator compute its own quality factor Q , this stability has a price, which is the loss of social optimality.

6 Incentive compatibility analysis

As we have seen in Sect. 2.1.3, incentive-compatibility is defined as the property of a negotiation mechanism which makes telling the truth the best strategy for any agent, assuming the rest of the agents also tell the truth. Though there are negotiation models where incentive compatibility can be proved analytically [11], these proofs are difficult to derive in the nonlinear domain. This is specially true for heuristics approaches with a great degree of variability, such as the model we are dealing with. In these cases, experimental evaluations may be conducted to assess the possible influence of insincere revelation of information over the stability of the negotiations. This is the approach we have taken to study incentive-compatibility in our model.

6.1 Experimental settings

Incentive compatibility analysis is oriented to evaluate the possibility for negotiating agents to manipulate the negotiation to their own benefit by means of revealing insincere information. In the negotiation model we are dealing with, information revealed to the mediator is the set of agents' bids. These bids represent regions within the solution space. Each offer has an associated utility value, a volume, and an associated quality factor value. Since bid volume is directly related to the region represented by the bid, it does not seem feasible to fake it, since it can be easily checked by the mediator. Quality factor may be faked, but since the mediator is very likely to recompute it using its own α parameter, this strategy is also harmless. Finally, agents may fake bid utility. Insincere information revelation about bid utility may generally occur in two ways: exaggerating upward or downward the utility values of *all* bids, or

exaggerating the utility values of *some* bids with respect to the others. Exaggerating all bids is not profitable with the proposed deal identification mechanisms, since bid selection at the mediator is performed independently for each agent. This means that the bids from different agents do not compete among each other to be selected as part of a solution. In contrast, the different bids of a single agent compete among themselves. Taking this into account, an agent could try to exaggerate the utility value of its preferred bids, thus trying to increase the probability of the mediator choosing those preferred bids to form deals. As far as social welfare is concerned, this is a problem if the set of exaggerated bids is small with respect to the total set of bids, since that would reduce the number of effective bids considered by the mediator, thus reducing deal probability.

To study the effect of utility exaggerations over the negotiations, we have conducted experiments comparing the utility obtained by an *insincere agent* with the utility obtained being sincere, assuming the rest of the agents are sincere. The behavior of the insincere agent is modeled by exaggerating the utility of a portion of the agent's highest utility bids. We have considered different degrees of exaggeration for the insincere agent.

- *Reference*: There are no insincere agents.
- *75%*: The insincere agent exaggerates 75% of its bids.
- *50%*: The insincere agent exaggerates half of its bids.
- *25%*: The insincere agent exaggerates one quarter of its bids.
- *12.5%*: The insincere agent exaggerates one eighth of its bids.

In all cases, exaggerated bids are the ones which yield better utility for the agents before exaggeration. Bid exaggeration is performed by multiplying the affected bids by a constant. The constant has been chosen to be higher than the average utility for agent bids, in order to make more likely that exaggeration could significantly impact the mediator's choice. In these experiments, the value for this constant is 10000. Again, experiments have been repeated for utility spaces with different values for the correlation length ψ .

6.2 Experimental results

Experiment results for SA-Q y MWIS-Q for six agents and six issues are shown, respectively, in Tables 7 and 8. Each table represents the median ratios between the utilities obtained by insincere and truthful agents. The results are statistically significant for $P < 0.05$. We can see that there are only significant gains for the insincere agents in medium complexity scenarios. In high-complexity scenarios, the presence of the insincere agent makes the negotiations fail, and thus there is no incentive to deviate from the socially optimal strategy. When utility space complexity decreases, we can see that an insincere agent may obtain gains over 40% for both SA-Q and MWIS-Q depending on the degree of exaggeration. Finally, for the less-complex scenarios, insincere revelation of information does not imply a significant gain in utility, since all agents achieve high utility values by being sincere.

Figure 12a and b show the box plots of the results for 100 runs of the experiments for SA-Q and MWIS-Q, in the most critical scenarios identified above (i.e. $\psi = 4.0$ for SA-Q and $\psi = 4.3$ for MWIS-Q). We can see a different evolution in the gain for the insincere agent as the degree of exaggeration varies. For SA-Q, this gain increases as the proportion of exaggerated bids decreases, which is reasonable taking into account that, if the mediator is successfully tricked into choosing bids only from the exaggerated set, the average utility of the bids in the set is higher (they are its better n bids). Exaggerating too much, however, can excessively reduce the selected bid set, thus impacting deal probability and making negotiations fail, which happens for a 12.5% degree of exaggeration. For MWIS-Q the maximum

Table 7 Incentive-compatibility analysis for SA-Q

ψ	Degree of exaggeration				
	Reference	75%	50%	25%	12.5%
2.8	0.9875	1.0061	–	–	–
3.1	0.9903	1.0107	–	–	–
4.0	0.9904	1.2708	1.4464	1.5071	–
4.3	0.9882	0.9662	1.0042	0.9727	0.9981
4.6	1.0015	1.0037	0.9858	0.9866	0.9974
5.9	1.0042	1.0107	1.0010	0.9840	1.0040

Table 8 Incentive-compatibility analysis for MWIS-Q

ψ	Degree of exaggeration				
	Reference	75%	50%	25%	12.5%
2.8	1.0022	0.9656	–	–	–
3.1	0.9783	0.9777	–	–	–
4.0	1.0035	1.0051	–	–	–
4.3	0.9763	1.1459	1.4785	1.3523	1.1614
4.6	0.9882	0.9463	0.9991	0.9672	0.9968
5.9	1.0091	1.0145	1.0054	0.9544	1.0139

gain is achieved for 50% degree of exaggeration, and further narrowing of the exaggerated bid set makes the gain for the insincere agent decrease, but it does not make negotiations fail. This is an effect of the higher correlation in the MWIS-Q selected scenario ($\psi = 4.3$), which makes deal probability higher. Finally, we can observe that exaggeration of the 75% of the bids has no significant effect, since most agent bids are included in the exaggerated set in this case. From these results we can conclude that there are incentives for the agents to behave insincerely in those scenarios, and therefore additional mechanisms should be introduced in the model to make it incentive-compatible.

6.3 Incentivizing sincere behavior in the auction-based negotiation protocol

As we have seen, the proposed model is prone to manipulations by means of *exaggerations* made by the agents, and there is an incentive for agents to behave insincerely. This is an undesirable property in a negotiation model, and may lead to further stability problems. Therefore, we seek for mechanisms which counter this effect, incentivizing sincere revelation of information. A possibility to achieve this is to normalize the utility values assigned by the agents to their bids, thus lowering the absolute differences in utility. We propose three different possibilities regarding utility normalization:

- *Normalization to maximum utility* : obtained by dividing each agent's bid utility by the maximum utility value issued by that agent:

$$u_n(b_i) = \frac{u(b_i)}{\max_{b_j \in B} u(b_j)}. \quad (6)$$

Using this normalization mechanism we can avoid the manipulation of the final deal by exaggerating upwards the utility values of the preferred offers. It does not prevent,

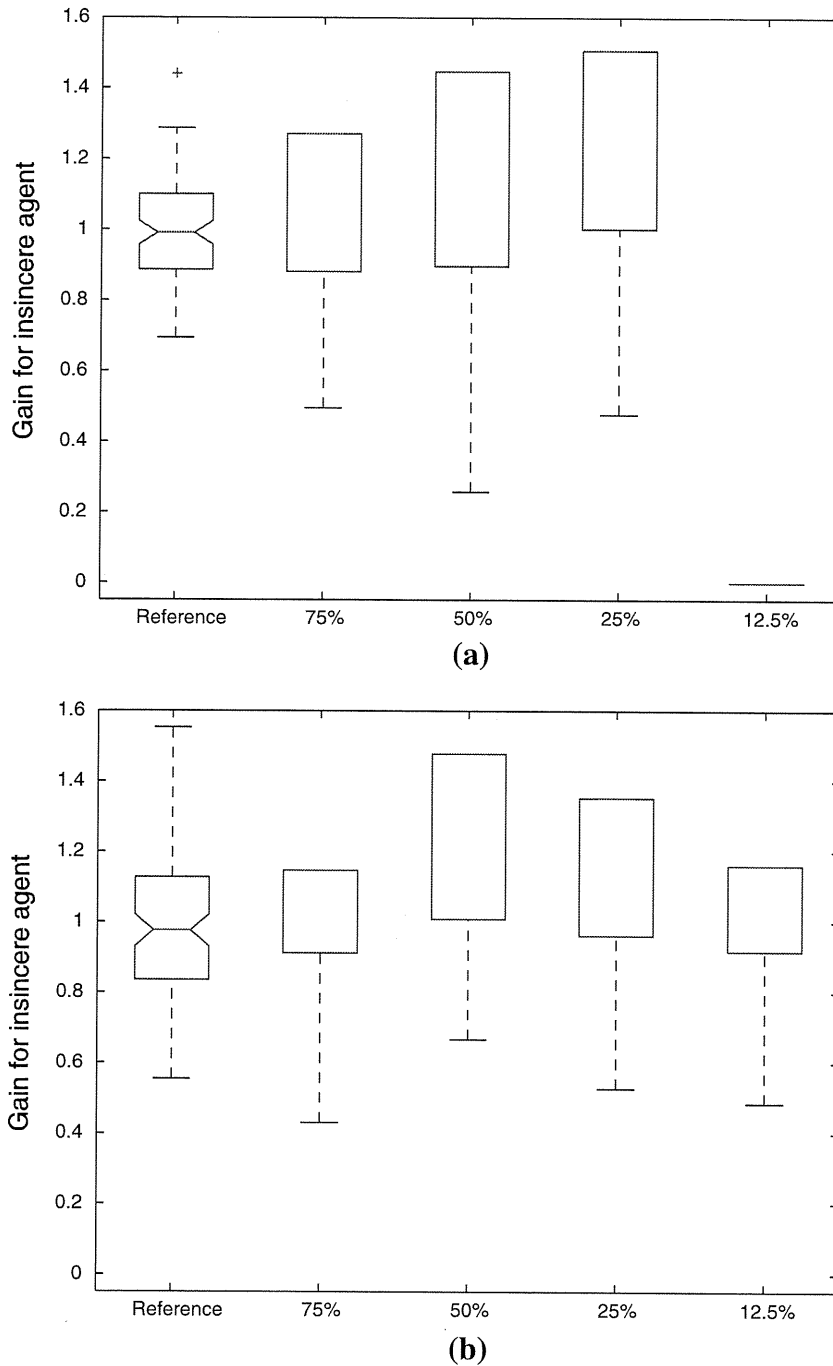


Fig. 12 Detail of the incentive-compatibility analysis for the most critical scenarios. **a** SA-Q $\psi = 4.0$, **b** MWIS-Q $\psi = 4.3$

- however, downward exaggerations, that is, to assign an extremely low value to the bids which are less profitable for the agent.
- *Bounded maximum-minimum normalization*: Attempts to prevent the manipulation of the negotiation model through upwards or backwards exaggerations. It is given by the expression

$$u_n(b_i) = u'_{\min} + \frac{u(b_i) - u_{\min}}{u_{\max} - u_{\min}} (u'_{\max} - u'_{\min}), \quad (7)$$

where $u_{\max} = \max_{b_j \in B} u(b_j)$, $u_{\min} = \min_{b_j \in B} u(b_j)$ and u'_{\min} and u'_{\max} are parameters chosen by the mediator. In this way, a utility mapping from the interval $[u_{\min}, u_{\max}]$ to the interval $[u'_{\min}, u'_{\max}]$ is performed for all bids, putting an upper bound $\frac{u'_{\max}}{u'_{\min}}$ to the ratio between the utilities of an agent's bids.

- *Ordinal normalization*: obtained by ordering the different bids of an agent according to their utility or quality factor, and mapping this order to a monotonically increasing succession of utility values, regardless of the original utility values. For instance, if B is the set of bids for an agent, in ascendent order of utility, and taking the arithmetic succession $s = \{1, 2, \dots, n_B\}$ as the mapping function, the normalized bid utility values would be of the form

$$u_n(b_i) = s_i = i.$$

Our hypothesis is that using these normalization methods may positively contribute to the incentive-compatibility of the model. To evaluate the effect of the proposed mechanisms over the incentive-compatibility of the model we have repeated the experiments performed above for the different normalization mechanisms proposed:


1. *Reference*: Utility values are not normalized.
2. *Umax*: Mediator uses *normalization to maximum utility* (Eq. 6).
3. *Bounded*: Mediator uses *bounded maximum-minimum normalization* (Eq. 7).
4. *Ordinal*: Mediator uses *ordinal normalization*.

Figure 13 a and b show the box plots of the results for 100 runs of the experiments for SA-Q and MWIS-Q in the most critical scenarios identified above, that is, $\psi = 4.0$ with a 25% degree of exaggeration for SA-Q and $\psi = 4.0$ with a 50% degree of exaggeration for SA-Q. We can see similar trends for both cases. Though all proposed normalization techniques reduce the incentive for the insincere agent to exaggerate, only bounded maximum-minimum normalization makes the expected gain for the insincere agent negligible, thus effectively removing the incentive to exaggerate, improving incentive-compatibility of the model.


7 Concluding remarks

Situations of high price of anarchy, which imply that individual rationality drives the agents towards strategies which yield low individual and social welfares, should be avoided 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 we proposed in refs. [31,55]. This strategy analysis has started studying the existence of individual and social optimal strategy profiles. This has revealed the existence of an individual optimal 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



(a)



(b)

Fig. 13 Effect of the proposed normalization mechanisms for the most critical scenarios. **a** SA-Q, $\psi = 4.0$, 25% degree of exaggeration. **b** MWIS-Q, $\psi = 4.3$, 50% degree of exaggeration

benefit from using its individually optimal strategy, which raises stability concerns, leading the model to high expected price of anarchy values. 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, avoiding the situations of infinite expected price of anarchy.

Finally, incentive compatibility issues in the protocol have been analyzed, showing that the model may be manipulated by agents which exaggerate the utility values of a subset of their bids, achieving significant gains for the insincere agents in medium correlated scenarios. To solve this, a set of normalization techniques have been proposed in order to incentivize sincere behavior. Experiments have shown that, though all proposed techniques reduce the incentive for an agent to exaggerate its bids, only the proposed bounded maximum-minimum normalization mechanism effectively removes the expected gain for being insincere, thus making the model incentive-compatible.

Though the experimental analysis performed has proven the effectiveness of the stability and incentive-compatibility mechanisms proposed, there is still plenty of research to be done in this area. We are interested on extending the strategy analysis presented in this work to an iterative version of the studied negotiation protocol, which would allow the agents to refine their bids in successive iterations of the protocol. This would raise very interesting additional considerations regarding agent and mediator strategies, since it would allow to develop adaptive measures. For the negotiation agents, this would mean, for instance, to be able to acquire a reasonable belief about the other agents' strategies during the negotiation, and to adapt its own strategy accordingly. This would drastically change the strategy analysis, since it would have to be conducted in a similar manner to a Bayes-Nash problem. The different results of the strategy analysis would probably impact the mechanisms needed at the mediator, and even more taking into account that the mediator could also take advantage of adaptive measures, trying to deduce agent strategies during the negotiation process, and 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 length 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.

Acknowledgements This work has been partially developed in the framework of the ITEA-2 project 2008005, "Do-it-Yourself Smart Experiences", and partially supported by the Spanish Ministry of Education and Science grant TIN2008-06739-C04-04, "T2C2".

A Appendix: Deduction of the expressions used in the probabilistic analysis

This section deduces these expression used in Sect. 4.1 for the probabilistic analysis of the auction based negotiation model. For ease of understanding, the deduction of the expressions is presented in a progressive manner. First of all, deal probability is calculated for an exchange between two agents of an *elemental* bid (a single unitary bid for each agent, for a single issue), and then it is shown how the expression varies when the number of issues and agents increase. Then, the resulting expression is generalized for an arbitrary number of bids per agent. Finally, given the expression for deal probability, expressions for expected utility and expected deal utility (defined as the expected utility conditioned to the event of a successful deal) are determined.

A.1 Deal probability

Considering the negotiation protocol described in Sect. 3.2, the probability of finding a deal is given by the probability of finding a common intersection of at least one bid of each agent. The simplest scenario we can devise is a bilateral, single issue negotiation where each agent makes a single, elemental bid, that is, a bid that represents a single point in the solution space.

Let a y b be the negotiating agents, and let $x^a, x^b \in D$ be their respective offers in a finite domain D with cardinality $|D|$. The probability P_{solution} of a deal or solution to the negotiation problem in this case is given by the probability of the coincidence of both bids. In this way,

$$\begin{aligned}
 P_{\text{solution}} &= \bigcup_{x \in D} p \left[(x^a = x) \cap (x^b = x) \right] = \underbrace{\sum_{x \in D} p \left[(x^a = x) \cap (x^b = x) \right]}_{x^j=x \text{ events are disjoint}} \\
 &= \underbrace{\sum_{x \in D} p(x^a = x) p(x^b = x)}_{x^a \text{ and } x^b \text{ are independent}} = \sum_{x \in D} \frac{1}{|D|} \frac{1}{|D|} = \frac{|D|}{|D|^2} = \frac{1}{|D|}, \tag{8}
 \end{aligned}$$

where we have assumed as the probability that a bid has a given value $p(x^a = x) = \frac{1}{|D|}$, which corresponds to a uniform bid distribution, for a maximum uncertainty scenario.

Extending the previous expression to a bilateral negotiation about n issues is straightforward. Again, let us consider the simplest case of a single elemental bid per agent. In this case, each bid will represent a point in an n -dimensional solution space, and the deal probability will be given by the probability of *all* issue values corresponding to one agent’s bid matching the respective values of the issues corresponding to the other agent’s bid.

Let a y b be the negotiating agents, and let $x^a, x^b \in D$ be their respective offers, such that the bid issued by agent j is given by $\bar{x}^j = \{x_i^j \mid i \in 1, \dots, n\}$, and such that $x_i^j \in D \forall i, j$. The probability of a deal or solution to the negotiation problem in this case is given by the expression

$$\begin{aligned}
 P_{\text{solution}} &= \bigcap_{1 \leq i \leq n} \left\{ \bigcup_{x \in D} p \left[(x_i^a = x) \cap (x_i^b = x) \right] \right\} \\
 &= \underbrace{\prod_{1 \leq i \leq n} \left\{ \bigcup_{x \in D} p \left[(x_i^a = x) \cap (x_i^b = x) \right] \right\}}_{\text{issue matches are independent events}} = \prod_{1 \leq i \leq n} \frac{1}{|D|} = \frac{1}{|D|^n}. \tag{9}
 \end{aligned}$$

In a similar way, this expression may be generalized to the case of n_a agents, taking into account that deal probability in this case is given by the probability of a match between the respective values for *all* issues of *all* agents’ bids, and that each agent bid is independent from the others’. In this way, the expression for the probability of finding a solution or deal in this case will be the following:

$$\begin{aligned}
 P_{\text{solution}} &= \bigcap_{1 \leq i \leq n} \left\{ \bigcup_{x \in D} p \left[\bigcap_{1 \leq j \leq n_a} (x_i^j = x) \right] \right\} \\
 &= \bigcap_{1 \leq i \leq n} \left\{ \bigcup_{x \in D} \left[\prod_{1 \leq j \leq n_a} p(x_i^j = x) \right] \right\} \\
 &= \bigcap_{1 \leq i \leq n} \left\{ \sum_{x \in D} \left[\prod_{1 \leq j \leq n_a} p(x_i^j = x) \right] \right\}
 \end{aligned}$$

$$\begin{aligned}
 &= \bigcap_{1 \leq i \leq n} \left\{ \sum_{x \in D} \left[\prod_{1 \leq j \leq n_a} \frac{1}{|D|} \right] \right\} \\
 &= \bigcap_{1 \leq i \leq n} \left\{ \frac{|D|}{|D|^{n_a}} \right\} = \prod_{1 \leq i \leq n} \left\{ \frac{1}{|D|^{n_a-1}} \right\} \\
 &= \frac{1}{|D|^{n(n_a-1)}}. \tag{10}
 \end{aligned}$$

So far we have considered only single, elemental bids, that is, each agent issued a single bid representing a single point in the solution space. This assumption allowed us to ensure deal events were disjoint (there was only a possible deal), which allowed to compute probabilistic unions as sums of probabilities. Generalizing to the case of multiple bids makes multiple points of agreement possible, and thus makes necessary to take into account possible intersection among deal events to compute probabilistic unions.

Given a set of N events E_1, \dots, E_N , with known probabilities $p(E_i)$, and not necessarily disjoint, the probability of the union $\bigcup_{i=1}^N E_i$ is given by

$$p\left(\bigcup_{i=1}^N E_i\right) = 1 - p\left(\bigcap_{j=1}^N \overline{E}_j\right).$$

If the events are independent and equiprobable, we have that $p(E_i) = p$, $p(\overline{E}_i) = 1 - p$, and the probability of the intersection above is given by $p(\bigcap_N \overline{E}_i) = (1 - p)^N$. In this case, we can see that the above expression leads to the following:

$$\begin{aligned}
 p\left(\bigcup_{i=1}^N E_i\right) &= 1 - (1 - p)^N \\
 &= 1 - \sum_{j=0}^N \binom{N}{j} 1^{N-j} (-p)^j \\
 &= 1 - \sum_{j=0}^N \binom{N}{j} (-1)^j p^j \\
 &= \sum_{j=1}^N (-1)^{j+1} \binom{N}{j} p^j. \tag{11}
 \end{aligned}$$

This result can be used to generalize the expression for deal probability obtained in the previous section to the case of multiple offers. Let us consider again a set of n_a agents negotiating about n issues. In this case we will consider that each agent k sends n_{bp}^k elemental bids. We consider elemental bids without loss of generality, since any other kind of bids (e.g. hyper-rectangles) can be decomposed to elemental bids. There may be overlaps between the different bids of an agent (i.e. they may or may not be disjoint). The probability P_{solution} that there is a solution or deal to the negotiation problem will be given by the probability that *at least one* of the possible combinations of bids from the different agents results in a deal. If each agent k issues n_{bp}^k bids there are $\prod n_{bp}^k$ possible combinations of one offer of each agent. The event C_l denotes the fact that the combination l results in a deal. The different events C_l

are equiprobable, and their probability is given by Eq. 10, reproduced here for convenience:

$$p(C_l) = \frac{1}{|D|^{n(n^a-1)}}.$$

Taking this into account, and using Eq. 11 for the computation of the probability of a union of equiprobable events, deal probability for the set of bids is given by the expression

$$\begin{aligned} P_{\text{solution}} &= p\left(\bigcup_{l=1}^{\prod n_{bp}^k} C_l\right) = \sum_{j=1}^{\prod n_{bp}^k} (-1)^{j+1} \binom{\prod n_{bp}^k}{j} p(C_l)^j \\ &= \sum_{j=1}^{\prod n_{bp}^k} (-1)^{j+1} \binom{\prod n_{bp}^k}{j} \left(\frac{1}{|D|^{n(n^a-1)}}\right)^j. \end{aligned} \tag{12}$$

A.2 Expected utility and expected deal utility

Once deal probability has been determined, it is easy to compute expected utility. By definition, the expected value of a random variable X which takes values from a domain D is computed as the sum $\sum_{x \in D} x \cdot p(X = x)$ of the products of each possible value for the variable and the respective probability that the variable takes each value. For the case of the expected utility for an agent, the possible values of the variable are the utility values associated to the different bids, and the probability that the variable takes each value is the probability that each bid results in a deal. To compute this probability, we have to take into account that each elemental bid \bar{x}_i^j of an agent j may be part of $\prod_{k \neq j} n_{bp}^k$ events $C(\bar{x}_i^j)_l$, representing the fact that the different combinations of this bid with the different elemental bids of the rest of the agents may result in a deal. In this way, the deal probability for a given elemental bid \bar{x}_i^j is given by

$$p(\bar{x}_i^j) = p\left(\bigcup_{l=1}^{\prod_{k \neq j} n_{bp}^k} C(\bar{x}_i^j)_l\right) = \sum_{j=1}^{\prod_{k \neq j} n_{bp}^k} (-1)^{j+1} \binom{\prod_{k \neq j} n_{bp}^k}{j} \left(\frac{1}{|D|^{n(n^a-1)}}\right)^j.$$

From this expression, the expected utility for an agent j is computed as follows:

$$\begin{aligned} E[u^j] &= \sum_{i=1}^{n_{bp}^j} u(\bar{x}_i^j) p(\bar{x}_i^j) \\ &= \sum_{i=1}^{n_{bp}^j} \left[u(\bar{x}_i^j) \sum_{j=1}^{\prod_{k \neq j} n_{bp}^k} (-1)^{j+1} \binom{\prod_{k \neq j} n_{bp}^k}{j} \left(\frac{1}{|D|^{n(n^a-1)}}\right)^j \right] \\ &= \left[\sum_{i=1}^{n_{bp}^j} u(\bar{x}_i^j) \right] \left[\sum_{j=1}^{\prod_{k \neq j} n_{bp}^k} (-1)^{j+1} \binom{\prod_{k \neq j} n_{bp}^k}{j} \left(\frac{1}{|D|^{n(n^a-1)}}\right)^j \right], \end{aligned} \tag{13}$$

where $\sum_{i=1}^{n_{bp}^j} u(\bar{x}_i^j)$ is the sum of the utilities of all points issued as bids by the agent. For the case of non-elemental bids, we consider each agent j issues n_b^j bids. Each bid m of the

agent represents an *iso-surface* of the agent's preference space (e.g, an hyperrectangle), and thus may be decomposed in v_m^j elemental bids of the same utility u_m^j , where v_m^j is the *volume* of the iso-surface represented by the bid m . In this case, we can establish the equivalence $\sum_{i=1}^{n_{bp}^j} u(\bar{x}_i^j) = \sum_{m=1}^{n_b^k} u_m^k \cdot v_m^k$, and the expression for the expected utility results as follows

$$E[u^j] = \left[\sum_{m=1}^{n_b^k} u_m^k \cdot v_m^k \right] \left[\sum_{j=1}^{\prod_{k \neq j} n_{bp}^k} (-1)^{j+1} \binom{\prod_{k \neq j} n_{bp}^k}{j} \left(\frac{1}{|D|^{n(n^a-1)}} \right)^j \right],$$

which is the expression we saw for Eq. 2.

Finally, expected deal utility for an agent may be obtained easily, since it only depends on the utility distribution within the set of bids issued by the agent. Assuming a deal have been reached, the probability for each elemental bid to be part of the deal will be $p(\bar{x}_i^j | deal) = \frac{1}{n_{bp}^j}$, assuming the different elemental bids are equiprobable (maximum uncertainty scenario). Taking this into account, expected deal utility is given by

$$E[u^j | deal] = \sum_{i=1}^{n_{bp}^j} u(\bar{x}_i^j) p(\bar{x}_i^j | deal) = \frac{1}{n_{bp}^j} \sum_{i=1}^{n_{bp}^j} u(\bar{x}_i^j),$$

which, for hyperrectangular bids, takes the form we saw in Eq. 3:

$$E[u^j | deal] = \frac{\sum_{m=1}^{n_b^j} u_m^j \cdot v_m^j}{n_{bp}^j}. \quad (14)$$

References

1. An, B., Lesser, V., & Sim, K. (2010). Strategic agents for multi-resource negotiation. *Autonomous Agents and Multi-Agent Systems*, 1–40. doi:10.1007/s10458-010-9137-2.
2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., & Roughgarden, T. (2004). The price of stability for network design with fair cost allocation. In *FOCS '04: Proceedings of the 45th annual IEEE symposium on foundations of computer science* (pp. 295–304). Washington, DC: IEEE Computer Society. doi:10.1109/FOCS.2004.68.
3. Bayati, M., Shah, D., & Sharma, M. (2008). Max-product for maximum weight matching: Convergence, correctness, and lp duality. *IEEE Transactions on Information Theory*, 54(3), 1241–1251.
4. Beer, M., D'Inverno, M., Luck, M., Jennings, N., Preist, C., & Schroeder, M. (1999). Negotiation in multi-agent systems. *Knowledge Engineering Review*, 14(3), 285–290.
5. Bichler, M., Kaukal, M., & Segev, A. (1999). Multi-attribute auctions for electronic procurement. In *Proceedings of the first IBM IAC. Workshop on internet based negotiation technologies, Yorktown Heights, NY*.
6. Buttner, R. (2006). A classification structure for automated negotiations. In *WI-IATW '06: Proceedings of the 2006 IEEE/WIC/ACM international conference on web intelligence and intelligent agent technology* (pp. 523–530). Washington, DC. doi:10.1109/WI-IATW.2006.7
7. Chevaleyre, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-aguilar, J. A., & Sousa, P. (2006). Issues in multiagent resource allocation. *Information*, 30, 2006.
8. Chevaleyre, Y., Endriss, U., Estivie, S., & Maudet, N. (2008). Multiagent resource allocation in k-additive domains: Preference representation and complexity. *Annals of Operations Research*, 163, 49–62. doi:10.1007/s10479-008-0335-0.
9. Choi, S. P. M., Liu, J., Chan, S. P. (2001). A genetic agent-based negotiation system. *Computer Networks*, 37(2), 195–204 (electronic business systems)

10. Chou, T. C., Fu, L. C., & Liu, K. P. (2007). E-negotiation of dependent multiple issues by using a joint search strategy. In *ICRA'07: IEEE international conference on robotics and automation* (pp. 1298–1303)
11. Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, *11*(1), 17–33. doi:10.1007/BF01726210.
12. Ehtamo, H., Hamalainen, R. P., Heiskanen, P., Teich, J., Verkama, M., & Zionts, S. (1999). Generating pareto solutions in a two-party setting: Constraint proposal methods. *Management Science*, *45*(12), 1697–1709.
13. Faratin, P., Sierra, C., & Jennings, N. R. (1998). Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, *24*(3–4), 159–182.
14. Faratin, P., Sierra, C., & Jennings, N. R. (2002). Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, *142*(2), 205–237.
15. Fatima, S., Wooldridge, M. J., & Jennings, N. R. (2004). An agenda based framework for multi-issues negotiation. *Artificial Intelligent Journal*, *152*(1), 1–45.
16. Fatima, S., Wooldridge, M., & Jennings, N. R. (2006). Multi-issue negotiation with deadlines. *Journal of Artificial Intelligence Research* *27*, 381–417. <http://eprints.ecs.soton.ac.uk/13079/>.
17. Fatima, S., Wooldridge, M., & Jennings, N. R. (2007). Approximate and online multi-issue negotiation. In *AAMAS '07: Proceedings of the 6th international joint conference on autonomous agents and multiagent systems* (pp. 1–8). ACM, New York. doi:10.1145/1329125.1329315
18. Fatima, S., Wooldridge, M., & Jennings, N. R. (2009). An analysis of feasible solutions for multi-issue negotiation involving nonlinear utility functions. In *AAMAS '09: Proceedings of the 8th international conference on autonomous agents and multiagent systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
19. Fristrup, P., & Kleiding, H. (1989). A note on asymptotical strategy-proofness. *Economics Letters*, *31*(4), 307–312. <http://ideas.repec.org/a/eee/econlet/v31y1989i4p307-312.html>.
20. Gatti, N., & Amigoni, F. (2005). An approximate pareto optimal cooperative negotiation model for multiple. In *IAT '05: Proceedings of the IEEE/WIC/ACM international conference on intelligent agent technology* (pp. 565–571). Washington, DC: IEEE Computer Society. doi:10.1109/IAT.2005.40
21. Giovannucci, A., Cerquides, J., Endriss, U., & Rodríguez-Aguilar, J. (2010). A graphical formalism for mixed multi-unit combinatorial auctions. *Autonomous Agents and Multi-Agent Systems*, *20*, 342–368. doi:10.1007/s10458-009-9085-x.
22. Grabisch, M. (1997). k-Order additive discrete fuzzy measures and their representation. *Fuzzy Sets and Systems*, *92*(2), 167–189. doi:10.1016/S0165-0114(97)00168-1.
23. Guttman, R. H., Moukas, A. G., & Maes, P. (1998). Agent-mediated electronic commerce: A survey. *The Knowledge Engineering Review*, *13*(2), 147–159.
24. Harsanyi, J. C. (2004). Games with incomplete information played by bayesian players. *Management Science*, *50*(12 Supplement), 1804–1817.
25. He, M., Jennings, N. R., & Leung, H. F. (2003). On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering*, *15*(4), 985–1003.
26. Hindriks, K., Jonker, C., & Tykhonov, D. (2006). Eliminating interdependencies between issues for multi-issue negotiation. In *Cooperative information agents X, lecture notes in computer science*, vol 4149, (pp. 301–316). Berlin: Springer.
27. Hordijk, W. (1995). *A measure of landscapes*. Working Papers 95-05-049, Santa Fe Institute. <http://ideas.repec.org/p/wop/safiw/95-05-049.html>
28. Horn, J., & Goldberg, D. E. (1994). Genetic algorithm difficulty and the modality of fitness landscapes. In *Foundations of genetic algorithms* (Vol. 3, pp. 243–269). Los Altos: Morgan Kaufmann.
29. Hunsberger, L., & Grosz B. J. (2000). A combinatorial auction for collaborative planning. Multi-agent systems. In *International Conference on 0:0151*. doi:10.1109/ICMAS.2000.858447.
30. Ito, T., Klein, M., & Hattori, H. (2007). Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces. In *Proceedings of the 20th international joint conference on artificial intelligence (IJCAI07)* (pp. 1347–1352).
31. Ito, T., Klein, M., & Hattori, H. (2008). A multi-issue negotiation protocol among agents with nonlinear utility functions. *Journal of Multiagent and Grid Systems*, *4*(1), 67–83.
32. Jennings, N. R. (2001). An agent-based approach for building complex software systems. *Communications of the ACM*, *44*(4), 35–41.
33. Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Sierra, C., & Wooldridge, M. (2001). Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, *10*(2), 199–215.

34. Jonker, C., Robu, V., & Treur, J. (2007). An agent architecture for multi-attribute negotiation using incomplete preference information. *Autonomous Agents and Multi-Agent Systems*, 15, 221–252. doi:10.1007/s10458-006-9009-y.
35. Kalai, E. (1977). Nonsymmetric nash solutions and replications of 2-person bargaining. *International Journal of Game Theory*, 6(3), 129–133. doi:10.1007/BF01774658.
36. Keeney, R., & Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value trade-offs*. New York: Wiley.
37. Kersten, G. E., & Noronha, S. J. (1998). Rational agents, contract curves, and inefficient compromises. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 28(3), 326–338.
38. Klein, M., Faratin, P., & Bar-Yam, Y. (2002). Using an annealing mediator to solve the prisoners' dilemma in the negotiation of complex contracts. In *AAMAS '02: Revised papers from the workshop on agent mediated electronic commerce on agent-mediated electronic commerce IV, designing mechanisms and systems* (pp. 194–202). London: Springer-Verlag.
39. Klein, M., Faratin, P., Sayama, H., & Bar-Yam, Y. (2003). Protocols for negotiating complex contracts. *IEEE Intelligent Systems*, 18(6), 32–38.
40. Kraus, S. (2001a). Automated negotiation and decision making in multiagent environments. In *Multi-agents systems and applications* (pp. 150–172). New York, NY: Springer-Verlag Inc.
41. Kraus, S. (2001). *Strategic negotiation in multiagent environments*. Cambridge, MA: MIT Press.
42. Kraus, S., Sycara, K., & Evenchick, A. (1998). Reaching agreements through argumentation: A logical model and implementation. *Artificial Intelligence*, 1–2, 1–69.
43. Lai, G., Li, C., Sycara, K., & Giampapa, J. (2004). Literature review on multiattribute negotiations. In *Technical Report CMU-RI-TR-04-66*. Pittsburgh: Robotics Institute, Carnegie Mellon University.
44. Lai, G., Li, C., & Sycara, K. (2006). Efficient multi-attribute negotiation with incomplete information. *Group Decision and Negotiation*, 15(5), 511–528. doi:10.1007/s10726-006-9041-y.
45. Lau, R. Y., Tang, M., & Wong, O. (2004). Towards genetically optimised responsive negotiation agents. In I. C. Society (Ed.) *Proceedings of the IEEE/WIC/ACM international conference on intelligent agent technology (IAT'04)* (pp. 295–301). Beijing: IEEE Computer Society.
46. Leme, R. P., & Tardos, E. (2010). Pure and bayes-nash price of anarchy for generalized second price auction. In I. press (Ed.) *Proceedings of the 51st annual IEEE symposium on foundations of computer science (FOCS 2010)*.
47. Lin, M. W. (2004). Modeling agent negotiation via fuzzy constraints in e-business. *Computational Intelligence*, 20, 624–642.
48. Lopez-Carmona, M. A. (2006). *Estrategias de negociación automática basadas en restricciones difusas sobre sistemas multiagente*. PhD thesis, Universidad de Alcalá.
49. Lopez-Carmona, M. A., & Velasco, J. R. (2006). An expressive approach to fuzzy constraint based agent purchase negotiation. In *Proceedings of the international joint conference on autonomous agents and multi-agent systems (AAMAS-2006)* (pp. 429–431). Japan: Hakodate
50. Lopez-Carmona, M. A., Velasco, J. R., & Marsa-Maestre, I. (2007). The agents' attitudes in fuzzy constraint based automated purchase negotiations. In *Multi-agent systems and applications V, lecture notes in artificial intelligence* (Vol. 4696, pp. 246–255). Berlin: Springer Verlag.
51. Luo, X., Jennings, N. R., Shadbolt, N., & Ho-Fung-Leung, & Lee, J. H. M. (2003a). A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments. *Artificial Intelligence*, 148(1–2), 53–102.
52. Luo, X., Lee, J. H., Leung, H. F., & Jennings, N. R. (2003b). Prioritised fuzzy constraint satisfaction problems: Axioms, instantiation and validation. *Fuzzy Sets and Systems*, 136(2), 151–188.
53. Manderick, B., de Weger, M., & Spiessens, P. (1991). The genetic algorithm and the structure of the fitness landscape. In *Proceedings of the fourth international conference on genetic algorithms*, San Diego, CA (pp. 1143–150).
54. Marsa-Maestre, I., Lopez-Carmona, M. A., Velasco, J. R., & de la Hoz, E. (2009a). Effective bidding and deal identification for negotiations in highly nonlinear scenarios. In *Proceedings of the VIII international conference on autonomous agents and multiagent systems (AAMAS 2009)* (pp. 1057–1064).
55. Marsa-Maestre, I., Lopez-Carmona, M. A., Velasco, J. R., Ito T., Fujita, K., & Klein, M. (2009b). Balancing utility and deal probability for negotiations in highly nonlinear utility spaces. In *Proceedings of the twenty-first international joint conference on artificial intelligence (IJCAI-09)* (pp. 214–219).
56. Marsa-Maestre, I., Lopez-Carmona, M. A., Velasco, J. R., de la Hoz, E. (2010). Avoiding the prisoner's dilemma in auction-based negotiations for highly rugged utility spaces. In *Proceedings of 9th international conference on autonomous agents and multiagent systems (AAMAS 2010)* (pp. 425–432).

57. Miller, B. L., & Goldberg, D. E. (1995). Genetic algorithms, tournament selection, and the effects of noise. *Complex Systems*, 9(3), 193–212.
58. Myerson, R. B. (1983). Mechanism design by an informed principal. *Econometrica*, 51(6), 1767–1797.
59. Nash, J. F. (1950). The bargaining problem. *Econometrica*, 18(2), 155–162.
60. Nash, J. F. (1953). Two-person cooperative games. *Econometrica*, 21(1), 128–140.
61. Nisan, N. (2006). Bidding languages. In P. Cramton, Y. Shoham, & R. Steinberg (Eds.), *Combinatorial auctions*. Cambridge, MA: MIT Press.
62. Osborne, M., & Rubinstein, A. (1990). *Bargaining and markets*. San Diego: The Academic Press.
63. Papadimitriou, C. (2001). Algorithms, games, and the internet. In: *STOC '01: Proceedings of the thirty-third annual ACM symposium on theory of computing* (pp. 749–753). New York, NY: ACM. doi:10.1145/380752.380883
64. Parkes, D. C., & Kalagnanam, J. (2005). Models for iterative multiattribute procurement auctions. *Management Science*, 51(3), 435–451. doi:10.1287/mnsc.1040.0340.
65. Poundstone, W. (1993). *Prisoner's dilemma*. New York: Anchor.
66. Rahwan, I., Ramchurn, S. D., Jennings, N. R., Mcburney, P., Parsons, S., & Sonenberg, L. (2003). Argumentation-based negotiation. *Knowledge Engineering Review*, 18(4), 343–375. doi:10.1017/S0269888904000098.
67. Raiffa, H. (1982). *The art and science of negotiation*. Cambridge, MA: Harvard University Press.
68. Reeves, D. M., & Wellman, M. P. (2004). Computing best-response strategies in infinite games of incomplete information. In *UAI '04: Proceedings of the 20th conference on Uncertainty in artificial intelligence* (pp. 470–478). Arlington, VA: AUAI Press.
69. Robu V., Somefun D. J. A., & La Poutré J. A. (2005). Modeling complex multi-issue negotiations using utility graphs. In *AAMAS '05: Proceedings of the fourth international joint conference on autonomous agents and multiagent systems* (pp. 280–287). New York, NY: ACM.
70. Ros, R., & Sierra, C. (2006). A negotiation meta strategy combining trade-off and concession moves. *Autonomous Agents and Multi-Agent Systems*, 12(2), 163–181. doi:10.1007/s10458-006-5837-z.
71. Rosenschein, J. S., & Zlotkin, G. (1994). *Rules of encounter*. Cambridge, MA: MIT Press.
72. Sakurai, Y., Yokoo, M., & Kamei, K. (2000). An efficient approximate algorithm for winner determination in combinatorial auctions. In *EC '00: Proceedings of the 2nd ACM conference on electronic commerce* (pp. 30–37). New York, NY: ACM. doi:10.1145/352871.352875
73. Sandholm, T. (2002). Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1–2), 1–54. doi:10.1016/S0004-3702(01)00159-X, <http://www.sciencedirect.com/science/article/B6TYF-4475Y1H-1/2/aafde63993109762200291d55309bf2e>
74. Sawaragi, Y., Nakayama, H., & Tanino, T. (1985). *Theory of multiobjective optimization*. New York: Academic Press.
75. Shi, C., Luo, J., & Lin, F. (2006) A multi-agent negotiation model applied in multi-objective optimization. In *Agent computing and multi-agent systems* (pp. 305–314).
76. Sierra, C. (2004). Agent-mediated electronic commerce. *Autonomous Agents and Multi-Agent Systems*, 9, 285–301.
77. Teich, J. E., Wallenius, H., Wallenius, J., & Zaitsev, A. (1999). A multiple unit auction algorithm: Some theory and a web implementation. *Electronic Markets*, 9(3), 199–205.
78. Teich, J. E., Wallenius, H., Wallenius, J., Zaitsev, A. (2006). A multi-attribute e-auction mechanism for procurement: Theoretical foundations. *European Journal of Operational Research*, 175(1):90–100. doi:10.1016/j.ejor.2005.04.023, <http://www.sciencedirect.com/science/article/B6VCT-4GH49SH-1/2/617b408205d1df0a41b488455fcfa127>
79. Tomassini, M., Vanneschi, L., Collard, P., & Clergue, M. (2005). A study of fitness distance correlation as a difficulty measure in genetic programming. *Evolutionary Computation*, 13(2), 213–239. doi:10.1162/1063656054088549.
80. Vassilev, V. K., Fogarty, T. C., & Miller, J. F. (2003). *Smoothness, ruggedness and neutrality of fitness landscapes: from theory to application* (pp. 3–44). New York, NY: Springer-Verlag, Inc.
81. Vetsikas, I., & Jennings, N. (2010). Bidding strategies for realistic multi-unit sealed-bid auctions. *Autonomous Agents and Multi-Agent Systems*, 21, 265–291. doi:10.1007/s10458-009-9109-6.
82. Volgenant, A. (2002). Solving some lexicographic multi-objective combinatorial problems. *European Journal of Operational Research*, 139(3):578–584. doi:10.1016/S0377-2217(01)00214-4, <http://www.sciencedirect.com/science/article/B6VCT-45DBB6C-9/2/d228ddb90c316574e86a4a8ee6278aba>
83. Weiss, G. (1999). *Multiagent Systems: A modern approach to distributed artificial intelligence*. Cambridge MA, USA: MIT Press.
84. Wierzbicki, A. P., Kruś, L., & Makowski, M. (1993). The role of multi-objective optimization in negotiation and mediation support. *Theory and Decision*, 34(3), 201–214. doi:10.1007/BF01075189.

85. Woodruff R. S. (1952). Confidence intervals for medians and other position measures. *Journal of the American Statistical Association*, 47(260), 635–646. <http://www.jstor.org/stable/2280781>
86. Xia, M., Stallaert, J., & Whinston, A. B. (2005). Solving the combinatorial double auction problem. *European Journal of Operational Research*, 164(1), 239–251. doi:10.1016/j.ejor.2003.11.018, <http://www.sciencedirect.com/science/article/B6VCT-4BN0J81-2/2/4a0255f0b884b4203ed8c3899fd84e58>
87. Yager, R. (2007). Multi-agent negotiation using linguistically expressed mediation rules. *Group Decision and Negotiation*, 16(1), 1–23.
88. Yokoo, M. (2001). *Distributed constraint satisfaction*. Berline: Springer Verlag.
89. Zhang, X., Lesser, V., & Abdallah, S. (2005). Efficient management of multi-linked negotiation based on a formalized model. *Autonomous Agents and Multi-Agent Systems*, 10, 165–205. doi:10.1007/s10458-004-6978-6.

大規模交渉問題における論点グループ数の調整に基づいた自動合意形成機構

藤田 桂 英^{†1,†2} 伊藤 孝 行^{†1,†2} マーク クライン^{†2}

マルチエージェントシステムの分野において、各論点が相互依存関係にある複数交渉問題が注目されている。各論点が相互依存関係にあると効用関数は非線形になる。しかし、既存の手法では線形の効用関数を主な対象としており、非線形の効用関数を持つ膨大な計算コストへの対応が不十分であった。特に、論点が増加した場合の交渉プロトコルのスケーラビリティは重要な問題である。本論文では、論点の依存関係に基づいて論点グループを決定し、合意形成を行う手法を対象としてスケーラビリティの向上を目指す。さらに、本手法を複数ラウンドに拡張しグループ数を調整することで、最適性の高い合意案を発見できることを示す。

An Automated Consensus Mechanism Based on Adjustment of Issue-groups for Multi-issue Negotiation

KATSUhide FUJITA,^{†1,†2} TAKAYUKI ITO^{†1,†2}
and MARK KLEIN^{†2}

Most real-world negotiation involves multiple interdependent issues, which makes an agent's utility functions nonlinear. Traditional negotiation mechanisms, which were designed for linear utilities, do not fare well in nonlinear contexts. One of the main challenges in developing effective nonlinear negotiation protocols is scalability; they can produce excessively high failure rates, when there are many issues, due to computational intractability. One reasonable approach to reducing computational cost, while maintaining good quality outcomes, is to decompose the utility space into several largely independent sub-spaces. In this paper, we propose a method for decomposing a utility space based on issue interdependencies. We expand our mechanism to be multi-round as adjusting the number of issue-groups. This multi-round protocol can find high quality solutions with high scalability.