

MIT Open Access Articles

Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces

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

Citation: Lopez-Carmona, Miguel A. et al. "Addressing Stability Issues in Mediated Complex Contract Negotiations for Constraint-Based, Non-Monotonic Utility Spaces." *Autonomous Agents and Multi-Agent Systems* 24.3 (2012): 485–535.

As Published: <http://dx.doi.org/10.1007/s10458-010-9159-9>

Publisher: Springer US

Persistent URL: <http://hdl.handle.net/1721.1/105117>

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

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



Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces

Miguel A. Lopez-Carmona · Ivan Marsa-Maestre ·
Mark Klein · Takayuki Ito

Published online: 4 December 2010
© The Author(s) 2010

Abstract Negotiating contracts with multiple interdependent issues may yield non-monotonic, highly uncorrelated preference spaces for the participating agents. These scenarios are specially challenging because the complexity of the agents' utility functions makes traditional negotiation mechanisms not applicable. There is a number of recent research lines addressing complex negotiations in uncorrelated utility spaces. However, most of them focus on overcoming the problems imposed by the complexity of the scenario, without analyzing

M. A. Lopez-Carmona and I. Marsa-Maestre are visiting scholars at Massachusetts Institute of Technology (MIT) from Universidad de Alcalá.

M. A. Lopez-Carmona (✉) · I. Marsa-Maestre
Computer Engineering Department, Edificio Politecnico, Universidad de Alcalá,
Ctra. N-II, Km 33.6, 28805, Alcalá de Henares, Madrid, Spain
e-mail: miguelangel.lopez@uah.es

I. Marsa-Maestre
e-mail: ivan.marsa@uah.es

M. A. Lopez-Carmona · I. Marsa-Maestre · M. Klein
Center for Collective Intelligence, MIT Sloan School of Management, Massachusetts Institute
of Technology (MIT), 5 Cambridge Center, NE25-749A, Cambridge, MA 02142, USA

M. Klein
e-mail: m_klein@mit.edu

T. Ito
Department of Computer Science and Engineering, Graduate School of Engineering,
Nagoya Institute of Technology, Gokiso, Showa-ku, Nagoya 466-8555, Japan
e-mail: ito.takayuki@nitech.ac.jp

T. Ito
Todai Policy Alternatives Research Institute, University of Tokyo, 7-3-1 Hongo,
Bunkyo-ku, Tokyo 113-0033, Japan

T. Ito
Precursory Research for Embryonic Science and Technology (PRESTO),
Japan Science and Technology Agency (JST), Tokyo, Japan

the potential consequences of the strategic behavior of the negotiating 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. Specially problematic are high *price of anarchy* situations, which imply that individual rationality drives the agents towards strategies which yield low individual and social welfares. In scenarios involving highly uncorrelated utility spaces, “low social welfare” usually means that the negotiations fail, and therefore high price of anarchy situations should be avoided in the negotiation mechanisms. In our previous work, we proposed an auction-based negotiation model designed for negotiations about complex contracts when highly uncorrelated, constraint-based utility spaces are involved. This paper performs a strategy analysis of this model, revealing that the approach raises stability concerns, leading to situations with a high (or even infinite) price of anarchy. 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. Finally, incentive-compatibility of the model is studied.

Keywords Automated multi-issue negotiation · Complex utility spaces · Strategy analysis

1 Introduction

Automated negotiation provides an important mechanism to reach agreements among distributed decision makers [4,42,43,71]. It has been extensively studied from the perspective of e-commerce [23,25,49,76], though it can be seen from a more general perspective as a paradigm to solve coordination and cooperation problems in complex systems [38,32], providing a mechanism for autonomous agents to reach agreements on, e.g., task allocation, resource sharing, or surplus division [15,37].

A variety of negotiation models have been proposed according to the many different parameters which may characterize a negotiation scenario [6,43]. We briefly review the key concepts about multi-attribute negotiation and the most relevant works in the field in Sect. 2.1. In the last years, there has been an increasing interest in complex negotiations [39]. Complexity of a negotiation scenario may depend on several factors, like the cardinality of the solution space, the number of negotiating agents, the number of issues under negotiation, the degree of interdependency between the issues, and structural properties of the preference landscape of the different agents, like ruggedness, modality or correlation length [80]. Specially challenging are those scenarios involving high cardinality solution spaces, since they tend to make exhaustive search in the solution space highly inefficient, and those involving highly rugged or highly uncorrelated utility spaces, since traditional negotiation approaches (mostly intended for linear or quasi-concave utility functions) cannot be applied to these scenarios. We briefly discuss utility space complexity and the techniques used to measure it in Sect. 2.2.

We can find some successful research works in the literature addressing negotiation in nonlinear utility spaces. [39] presented, as far as we are aware, the first negotiation protocol specific for complex preference spaces, based on using simulated annealing to progressively enhance an agreement between two agents. In [26], a different approach is taken, reducing the complexity of the agent’s preference space by using approximations of the agents’ utility functions where issue interdependency has been removed. [17] do not study the inherent complexity of agent preference spaces, but the complexity introduced in a negotiation when agent preferences change over time. We comprehensively review these and other related works in Sect. 2.3.

In our previous work [54], we proposed a mediated, auction-based protocol for nonlinear utility spaces generated using weighted constraints, such as the ones we may encounter when negotiating complex contracts with multiple, interdependent clauses [30]. We also proposed a set of decision mechanisms to generate bids at the negotiating agents and to identify feasible deals at the mediator once the bids from the negotiating agents have been received [54]. We briefly summarize the approach in Sect. 3. Experiments showed that these approaches achieve high effectiveness (measured as high optimality rates and low failure rates for the negotiations) in moderately rugged utility spaces.

In [55], we extended this work to address highly-rugged utility spaces. We proposed the use of a technique to balance utility and deal probability in the negotiation process, which we called *quality factor*. This quality factor is used to bias bid generation and deal identification taking into account the agents' attitudes (e.g. risk attitude, selfishness, willingness to cooperate). From the mechanisms we proposed to take into account quality factor in the negotiations, the most successful ones are detailed in Sect. 3.4. The experiments showed 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' attitudes, our previous experiments limited these attitudes to a somewhat "cooperative" environment, where all agents have the same, neutral attitude. In a real, competitive environment, we expect to have agents with different attitudes interacting. This raises the problem of agent strategic behavior, which is introduced in Sect. 4. What happens when risk averse agents interact with risk willing agents? Is there an individually optimal strategy? If so, does this individually optimal strategy lead to satisfying solutions, or is the approach prone to situations where individual rationality lead to solutions of low social value? Furthermore, since the complexity 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 individually optimal strategy and the socially optimal strategy for different utility space complexity levels. From the results of the analysis we conclude that the auction-based protocol, as described in [55], has stability problems, leading to situations resulting in high expected price of anarchy (Sect. 4).
- We propose a set of mechanisms intended to improve protocol stability. 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 (Sect. 5).
- We separately study a specific stability concern, incentive compatibility, related to the possibility of agents manipulating the protocol by means of insincere revelation of information (Sect. 6).

For each contribution, an experimental evaluation has been performed to validate our hypothesis and evaluate its effect. The experimental settings are described in Sects. 4.2, 5.2, 6.2 and 6.3, along with the discussion of the results obtained. Finally, the last section summarizes our conclusions and sheds light on some future research.

2 Complex negotiation scenarios

In the last years, there has been an increasing interest in complex negotiation scenarios, where agents negotiate about multiple, interdependent issues [39]. These scenarios are spe-

cially challenging, since issue interdependency yields nonlinear utility spaces, which make classic negotiation approaches not applicable [30]. In this section we first briefly review existing research on multi-attribute negotiation and outline the key components of any negotiation model. Then we discuss the most relevant works so far on the field of agent-based complex automated negotiations. Finally, some of the issues raised by complex negotiation scenarios, which are directly relevant to our research, are described.

2.1 Multi-attribute negotiation

Multi-attribute negotiation may be seen as an interaction between two or more agents with the goal of reaching an agreement about a range of issues which usually involves solving a conflict of interests between the agents. This kind of interaction has been widely studied in different research areas, such as game theory [71], distributed artificial intelligence [13] and economics [67]. Using a notation similar to that used in [71] and [88], we can formally define a *multi-attribute negotiation domain* as a tuple

$$\langle X, D, Ag, U \rangle$$

where

- $X = \{x_i | i = 1, \dots, n\}$ is a finite set of variables, called attributes or issues;
- $D = \{d_i | i = 1, \dots, n\}$ is a finite set of domains, such that each domain d_i represents the feasible values of the variable x_i ;
- $Ag = \{1, \dots, m\}$ is the set of negotiating agents, also assumed finite;
- $U = \{U^j | j = 1, \dots, m\}$, where $U^j : D \rightarrow \mathbb{R}$ represents the preference structure or utility for agent j .

Multi-attribute negotiation is seen as an important challenge for the multi-agent system research community [43], and there is a great variety of negotiation models and protocols intended to address different parts of this challenge. These models may be classified according to different criteria [6], such as their structure, the dynamics of the negotiation process, or the different constraints (e.g. deadlines, information availability...). According to the theoretical foundations of the negotiation models, we can find approaches based on game theory, heuristic approaches and argumentation-based approaches. Game theory approaches aim to find optimal solutions analytically, analyzing equilibrium conditions [59]. These models are mathematically sound and elegant, but their practical use in some negotiation scenarios is somewhat restricted due to the assumptions usually made: unlimited computation and memory resources, perfect rationality and complete information. In *heuristic* approaches, however, these assumptions are relaxed, and participants attempt to find an “approximately-optimal” under bound rationality using heuristic search and evaluation methods [12–14, 20, 31, 39, 44, 70]. In *argumentation* based negotiation, agents are given the ability to reason their positions, including a meta-information level which allows them to use promises, rewards, threats and other incentives [66].

Regardless of the theoretical approach involved, different authors agree that there are three key components in a negotiation model [16, 33, 41]:

- An *interaction protocol*, which defines the rules of encounter among the negotiating agents, including what kind of offer exchange is allowed and what kind of deals may be reached and how they are established.
- The *preference* sets of the different agents, which allow them to assess the different solutions in terms of gain or utility and to compare them.

- A set of *decision mechanisms* and *strategies*, which govern agents' decision making, allowing them to determine which shall be their next action for a given negotiation state.

2.1.1 Interaction protocols for negotiation

The most-widespread interaction protocol for negotiation is based on the exchange of offers and counter offers, which are expressed as an assignation of values to the different attributes. This kind of negotiation protocols are known as positional bargaining. In argumentation based negotiation, however, this exchange of offers also includes meta-information, in order to allow reasoning about the positions of the different agents. A particular protocol family for multi-lateral negotiations are *auction-based protocols*, where negotiating agents send their offers (also called *bids*) to a mediator, which then decides the winning deal [77]. Auction-based protocols allow to efficiently deal with one-to-many and many-to-many negotiations. Another important division regarding interaction protocols is between *one-shot* protocols and *iterative* protocols. In one-shot protocols, there is a single interaction step between the agents [59]. In iterative protocols, on the other hand, agents have the opportunity to refine their positions in successive protocol iterations [62].

2.1.2 Preference sets, utility functions and the use of constraints

From the decision theory perspective, preferences express the absolute or relative satisfaction for an individual about a particular choice among different options [36]. [7] classify agent preference structures in four broad families: binary, ordinal, cardinal and fuzzy preference structures. Among these families, cardinal preference structures are probably the most widely used in complex negotiations. In particular, it is usual to define agent preferences by means of utility functions.

Formally, for a given multi-attribute domain $\langle X, D, Ag, U \rangle$, the *utility function* for each agent $j \in Ag$ is defined as

$$U^j : D \rightarrow \mathbb{R},$$

assigning to each possible combination of values in X or *deal* $s = \{s_i | i = 1, \dots, n; s_i \in d_i\}$ a real number, which represents the utility that deal s yields for agent j .

The most basic form to represent a utility function is to make an enumeration of the points in the solution space which yield a non-zero utility value. In this way, an agent's utility function may be represented as a set of pairs $\langle s, u(s) \rangle | u(s) \neq 0$, where $u(s)$ is the utility of the solution s for the agent. It is easy to see that, though this representation for utility functions is fully expressive, its cardinality may grow greatly with the number of issues or with the cardinality of each issue's domain. Because of that, more succinct representations for utility functions are used in most cases. Examples of such representations which are widely used in the negotiation literature are linear-additive utility functions [14] or k-additive utility functions [22].

Another widely used way to represent preferences and utility functions is the use of constraints over the values of the attributes. There is a vast variety of multi-attribute negotiation models and approaches making use of constraints in different forms, from hard constraints to soft, probabilistic or fuzzy constraints [31,47,52]. There are several reasons which favor the use of constraints in negotiation models. First, they allow for efficient methods for preference elicitation. Moreover, constraints allow to express dependencies between the possible values of the different attributes. Finally, the use of constraints for offer expression allow to limit the




Fig. 1 Example of a nonlinear utility space defined by means of weighted constraints

region of the solution space which has to be explored in a given negotiation step. Reducing the region of the utility space under exploration according to the constraints exchanged by agents is a widely used technique in automated negotiation [48], since it makes the search for agreements a more efficient process than when using positional bargaining, specially in complex negotiation scenarios.

A particular case of constraint-based utility representation which has been used to model complex utility spaces for negotiation are *weighted constraints*. There is a utility value for each constraint, and the total utility is defined as the sum of the utilities of all satisfied constraints.

More formally, the utility space of the agents may be defined as a set of constraints $C = \{c_k | k = 1, \dots, l\}$. Each constraint c_k has an associated utility value $u(c_k)$. If we note as $s \in x(c_k)$ the fact that a given contract $s = \{s_i | i = 1, \dots, n\}$ is in the set of contracts that satisfy constraint c_k , an agent’s utility for contract s may be 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. Figure 1 shows an example of the kind of utility spaces which may be modeled using weighted constraints.

2.1.3 Agent strategies, mechanism stability and incentive-compatibility

In an automated negotiation, a strategy guides the decision making process of an agent throughout the different stages of the negotiation protocol [41]. The main challenge in an automated negotiation scenario as far as decision mechanisms are concerned is to design *rational* agents, able to choose an adequate negotiation strategy. In negotiations among self-ish agents, negotiation mechanisms must be designed in a way that makes them stable, understanding stability as the impossibility (or at least difficulty) of the strategic manipulation of the mechanisms. This means that the mechanisms should motivate the agents to

act in an adequate way, since if a rational, selfish agent may benefit from taking a strategy which is different to the one expected by the protocol, it will do so. This problem is closely related to the notion of *equilibrium* defined in game theory. In an equilibrium, each player of the game has adopted a strategy that they have no rational incentive to change (because it is the best alternative, given the circumstances). There are different equilibrium conditions which can be defined, like dominant strategies [40, 83], Nash equilibrium [60] or Bayes-Nash equilibrium [24].

Achieving stability in a negotiation mechanism does not guarantee to reach solutions maximizing social welfare. Therefore, stability must not be used as a single criterion to evaluate decision mechanisms, and social welfare should also be considered. An specially illustrative example is the *prisoner's dilemma* [65], which describes an scenario where Nash equilibrium yields low utility values for the agents involved. A more generic concept which is becoming widely used to characterize situations where individual rationality leads agents to results which yield low social welfares is the notion of *price of anarchy*. The price of anarchy was first introduced in [63] in the context of selfish routing, as a measure of loss of social efficiency due to selfish behavior. In the context of a problem of social welfare maximization, price of anarchy can be defined as follows:

Definition 1 *Price of anarchy*. [63] The *price of anarchy (PoA)* in a given game is defined as the ratio between the social welfare of the best possible outcome of the game and the social welfare of the worst Nash equilibrium in the game:

$$PoA = \frac{\max_{s \in S} sw(s)}{\min_{s \in S_{Nash}} sw(s)},$$

where S is the set of all possible outcomes of the game, $S_{Nash} \subseteq S$ is the set of all possible outcomes induced by a Nash equilibrium in the game, and $sw(s)$ is the social welfare of a given outcome s .

Defined in this way, price of anarchy gives an indication of the potential loss in a given game when individually rational agents are confronted. In situations where PoA is high, additional mechanisms which incentivize social behavior are desirable, in order to modify the equilibrium conditions of the game and reduce this value of PoA, thus improving the stability of the protocol. Stability, however, may also come at a price. Even when worst-case equilibria can be avoided, equilibrium conditions may lead to solutions which are distant to the social optimum (generally due to the fact that stability enhancing measures favor “fair” solutions against Pareto-optimal ones). To measure this, price of stability is introduced in an analogous manner:

Definition 2 *Price of stability*. [2] The *price of stability (PoS)* in a given game is defined as the ratio between the social welfare of the best possible outcome of the game and the social welfare of the best Nash equilibrium in the game:

$$PoS = \frac{\max_{s \in S} sw(s)}{\max_{s \in S_{Nash}} sw(s)},$$

where S is the set of all possible outcomes of the game, $S_{Nash} \subseteq S$ is the set of all possible outcomes induced by a Nash equilibrium in the game, and $sw(s)$ is the social welfare of a given outcome s .

Taking this into account, when mechanisms are introduced to reduce price of anarchy in a game, their impact over price of stability should also be evaluated.

Another threat to mechanism stability is strategic revelation of information. In incomplete information scenarios [34], since the agents' beliefs about the preferences of a given agent may influence the decision mechanisms they use, an agent may use as a strategy to lie about its own preferences in order to manipulate the decision mechanisms of the rest of the agents to its own benefit. This raises an additional concern to mechanism design [83].

It would be desirable to design protocols which are not prone to be manipulated through insincere revelation of information. *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 incentive-compatibility is usually independently studied, it is closely related to the notions of strategic equilibrium seen above. In particular, incentive-compatibility may be seen as the property of a negotiation model where, regarding the possibility of telling or not telling the truth, having all agents telling the truth is a Nash equilibrium. A more restrictive property is (*strategy-proofness*), which imposes truthful revelation of information to be a dominant strategy. This means that for any agent the best choice is to tell the truth regardless of the other agents' attitudes towards sincerity [19].

An example of an incentive-compatible protocol is the *Vickrey auction*. The Vickrey auctions are second-price, sealed, one shot auctions. In this kind of auction, that an agent i bids above its real utility value $u_i(s)$ is a bad strategy, since there is a chance that the second highest bid is also above that utility value, which would imply that the agent would have to pay for the product more than its value. Furthermore, as Vickrey auction is second price, bidding below the utility level $u_i(s)$ is also a bad strategy, since it reduces the chance to bid without any advantage, as the price the agent will have to pay for the product is not given by its bid, but by the second highest bid. Another incentive compatible mechanism is the Clarke tax method [11], where a tax is imposed to each agent once the negotiation has ended, and this tax makes each agent "pay" for the impact that its participation had over other agents' utilities, showing that, in this way, if an agent's false valuation changes the negotiation result, the utility obtained by that agent (after taxes are applied) is never higher than the utility it would have gained using truthful valuations [83].

2.2 Negotiation, optimization and complexity

Though there has been an increasing interest in complex negotiations in the last years, little efforts have been made to study complexity itself within negotiation (apart from computational complexity, which has been thoroughly studied in many scenarios). Therefore, if we want to be able to assess complexity in negotiations, we need to resort to other knowledge areas. One area where many authors have dealt with complexity characterization and measurement is optimization. In fact, negotiation scenarios and optimization problems are often closely related, since there are many similarities in the ways both problem families are defined and addressed. For example, negotiating agents are usually utility optimizers, and negotiation mechanisms are often evaluated in terms of their ability to reach Pareto-optimal solutions. In negotiation, Pareto-optimal solutions are those where payoff cannot be improved for any of the agents without decreasing the payoff for another agent. This concept of Pareto-efficiency is also sought in multi-objective optimization, trying to find solutions where no further gains can be achieved in one of the objectives without losing in another [74]. Multi-objective optimization has been widely used for negotiation support [84], and negotiation mechanisms have also been used to solve multiobjective optimization problems, usually by distributing the different objectives among negotiating agents [75]. Therefore, some of the concepts studied in multiobjective optimization may be used in negotiation, and vice versa.

In the context of a multi-attribute negotiation, complexity of a given scenario may depend at least on the number of issues, the level of interdependency between the preferences on the issues, the domain of the issues, the possibility of change over time of the negotiation context, the method used to describe preferences and the structural properties of the agent's utility spaces. In general, a large number of issues with a high interdependency and a large domain contribute to more complex preference spaces. If the negotiation context changes over time, complexity also increases. The method to describe preferences also has an influence in the complexity of the negotiation scenario. This is specially true when optimization techniques are used to find high utility regions within an agent's utility spaces, or to find deals among different agents. A constraint-based preference space, for instance, may present discontinuities which make gradient based optimizers not applicable, while differentiable utility functions contribute to a faster local optimization. Therefore, to study complexity in negotiation scenarios, we may find useful to characterize structural complexity of the agents' utility spaces, and to this end we may benefit from existing research on function characterization for optimization.

In this context, and more specifically in the field of optimization using evolutionary algorithms, structural complexity analysis plays a crucial role, since algorithm search capabilities are greatly impacted by some structural properties of the optimized function, which is usually known as *fitness landscape* in evolutionary computation.

An interesting detail about fitness landscapes is that they include the definition of a *neighborhood operator* ϕ , which expresses the probability that the search function (usually, a genetic algorithm) passes from one point in the landscape to another [27]. This operator is directly related to the search mechanism used and its parameters (e.g. simulated annealing temperature or mutation probability for genetic algorithms), which implies an important consequence: the complexity of a utility space may be different depending on the considered search algorithm and its parameters. This operator also defines the concept of *neighbor solutions* in the space, which in turn influences the definition of local optima (maxima and minima), and therefore the structural properties of a fitness landscape which are interesting regarding search complexity within the space, such as modality [28], ruggedness, smoothness and neutrality [80].

Once the properties which has an influence on the complexity of a fitness landscape or a solution space have been studied, techniques which allow to measure the complexity of a given space are needed. Most of the approaches we can find in the literature are based on the correlation between different samples of the fitness function f , like *fitness distance correlation* metrics [79] or stochastic models representing the correlation structure of the space [27]. A metric which is easy to compute in most scenarios and allows to make quantitative evaluations about the complexity of a fitness or utility landscape is *correlation length* or *correlation distance*. Correlation distance is defined as the minimum distance ψ which makes correlation fall below a given threshold (usually 0.5), which gives an idea of the distance we can move throughout the solution space while keeping a certain correlation between samples [53].

2.3 Related research on automated negotiation in complex utility spaces

Klein et al. [39] present, as far as we are aware, the first negotiation protocols specific for complex preference spaces. They propose a simulated annealing-based approach, a refined version based on a parity-maintaining annealing mediator, and an unmediated version of the negotiation protocol. Of great interest in this work are the positive results about the use of simulated annealing as a way to regulate agent decision making, along with the use of agent expressiveness to allow the mediator to improve its proposals. However, this expres-

siveness is somewhat limited, with only four possible valuations which allow the mediator to decide which contract to use as a parent for mutation, but not in which direction to mutate it. On the other hand, the performed experiments only consider the bilateral negotiation scenario, though authors claim that the multiparty generalization is simple. Finally, the family of negotiation protocols they propose are specific for binary issues and binary dependencies. Higher-order dependencies and continuous-valued issues, common in many real-world contexts, are known to generate more challenging utility landscapes which are not considered in their work.

Luo et al. [51] propose a fuzzy constraint based framework for multi-attribute negotiations. In this framework a buyer agent defines a set of fuzzy constraints to describe its preferences. The proposals of the buyer agent are a set of hard constraints which are extracted from the set of fuzzy constraints. The seller agent responds with an offer or with a relaxation request. The buyer then decides whether to accept or reject an offer, or to relax some constraints by priority from the lowest to highest. In Lopez-Carmona and Velasco [49], Lopez-Carmona et al. [50] an improvement to Luo's model is presented. They devise an expressive negotiation protocol where proposals include a valuation of the different constraints, and seller's responses may contain explicit relaxation requests. It means that a seller agent may suggest the specific relaxation of one or more constraints. The relaxation suggested by a seller agent is based on utility and viability criteria, which improves the negotiation process. Though these constraint-based works model discontinuous preference spaces, the operators used to compute utility and the utility spaces defined yield monotonic preference spaces, which are far from the complex preference spaces covered in our work.

Another interesting approach to solve the computational cost and complexity of negotiating interdependent issues is to simplify the negotiation space. Hindriks et al. [26] propose a weighted approximation technique to simplify the utility space. They show that for smooth utility functions the application of this technique results in an outcome that closely matches the outcome based on the original interdependent utility structure. The method is evaluated for a number of randomly generated utility spaces with interdependent issues. Experiments show that this approach can achieve reasonably good outcomes for utility spaces with simple dependencies. However, an approximation error that deviates negotiation outcomes from the optimal solutions cannot be avoided, and this error may become larger when the approximated utility functions become more complex. Authors acknowledge as a necessary future work to study which kind of functions can be approximated accurately enough using this mechanism. Another limitation of this approach is that it is necessary to estimate a region of utility space where the actual outcome is expected to be (i.e. it is assumed that the region is known a priori by the agents).

In Robu et al. [69] utility graphs are used to model issue interdependencies for binary-valued issues. Utility graphs are inspired by graph theory and probabilistic influence networks to derive efficient heuristics for non-mediated bilateral negotiations about multiple issues. The idea is to decompose highly non-linear utility functions in sub-utilities of clusters of inter-related items. They show how utility graphs can be used to model an opponent's preferences. In this approach agents need prior information about the maximal structure of the utility space to be explored. Authors argue that this prior information could be obtained through a history of past negotiations or the input of domain experts. However, our approach has the advantage that outcomes can be reached without any prior information and that it is not restricted to binary-valued issues.

There are several proposals which employ genetic algorithms to learn opponent's preferences according to the history of the counter-offers based upon stochastic approximation. In Choi et al. [9] a system based on genetic-algorithms for electronic business is proposed. In

this work the utility functions are restricted to take a product combination form (i.e. utility of an outcome is the product of the utility values of the different issues). The objective function used is based on the comparison of the changes of consecutive offers. Small changes of an issue suggest that this issue is more important. For each new population, the protocol enforces that the generated candidates cannot be better than the previous offer. Unlike other negotiation models based on genetic algorithms, this proposal adapts to the environment by dynamically modifying its mutation rate. Lau et al. [45] have also reported a negotiation mechanism for non-mediated automated negotiations based on genetic algorithms. The fitness function relies on three aspects: an agent's own preference, the distance of a candidate offer to the previous opponent's offer, and time pressure. In this work agents' preferences are quantified by a linear aggregation of the issue valuations. However, non-monotonic and discontinuous preference spaces are not explored. In Chou et al. [10] a genetic algorithm is proposed which is based on a joint elitism operation and a joint fitness operation. In the joint elitism operation an agent stores the latest offers received from the opponent. The joint fitness operation combines agent's own utility function and euclidean distance to the opponent's offer. In this work two different negotiation scenarios are considered. In the first one utility is defined as the weighted sum of the different issue values (i.e. issues are independent). The second scenario defines a utility function where there is a master issue and a set of slave issues. Utility is calculated as the weighted sum of the different issue values, but the weights of the slave and master issues change according to the value of the master issue.

In Yager [87] a mediated negotiation framework for multi-agent negotiation is presented. This framework involves a mediation step in which the individual preference functions are aggregated to obtain a group preference function. The main interest is focused on the implementation of the mediation rule where they allow a linguistic description of the rule using fuzzy logic. A notable feature of their approach is the inclusion of a mechanism rewarding the agents for being open to alternatives other than simply their most preferred. The negotiation space and utility values are assumed to be arbitrary (i.e. preferences can be non-monotonic). However, the set of possible solutions is defined a priori and is fixed. Moreover, the preference function needs to be provided to the mediation step in the negotiation process, and pareto-optimality is not considered. Instead, the stopping rule is considered, which determines when the rounds of mediation stop.

Fatima et al. [18] analyze bilateral multi-issue negotiation involving nonlinear utility functions. They consider the case where issues are divisible and there are time constraints in the form of deadlines and discounts. They show that it is possible to reach Pareto-optimal agreements by negotiating all the issues together, and that finding an equilibrium is not computationally easy if the agents' utility functions are nonlinear. In order to overcome this complexity they investigate two solutions: approximating nonlinear utilities with linear ones; and using a simultaneous procedure where the issues are discussed in parallel but independently of each other. This study shows that the equilibrium can be computed in polynomial time. An important part of this work is the complexity analysis and estimated approximation error analysis performed over the proposed approximated equilibrium strategies. Heuristic approaches have generally the drawback of the lack of a solid mathematical structure which guarantees their viability, which raises the need of an exhaustive experimental evaluation. An adequate complexity analysis and establishing a bound over the approximation error contribute to give heuristic approaches part of the technical soundness they usually lack. Among the limitations of the proposal, we can point out that this work is focused on symmetric agents where the preferences are distributed identically, and the utility functions are separable in nonlinear polynomials of a single variable. This somewhat limits the complexity of the preference space.

Finally, combinatorial auctions [21,29,72,73,82,86] can enable large-scale collective decision making in nonlinear domains, but only of a very limited type (i.e. negotiations consisting solely of resource allocation decisions). Multi-attribute auctions, wherein buyers advertise their utility functions, and sellers compete to offer the highest-utility bid [5,78,64] are also aimed at a fundamentally limited problem (a purchase negotiation with a single buyer) and require full revelation of preference information.

In summary, in the existing research nearly all the models which assume issue interdependency rely on monotonic utility spaces, binary valued issues, low-order dependencies, or a fixed set of defined a priori solutions. Simplification of the negotiation space has also been reported as a valid approach for simple utility functions, but it cannot be used with higher-order issue dependencies, which generate highly uncorrelated utility spaces. Therefore, new approaches are needed if automated negotiation is to be applied to settings involving non-monotonic, highly uncorrelated preference spaces.

3 An auction based approach for negotiations in highly uncorrelated, constraint based utility spaces

In this work we analyze agents' strategic behavior and mechanism stability for a mediated, auction-based negotiation approach we designed for highly uncorrelated, constraint based utility spaces [30,54]. To make such strategic analysis easier to understand, in this section we motivate and review the most relevant aspects of our negotiation model.


3.1 Negotiation domain and agent preference model

We explore the problem of negotiating complex contracts, which was first introduced by Klein et al. [39]. Contracts are defined as a set of issues or *clauses*, each of which may have a value. The aforementioned authors limited encounters to bilateral negotiations (i.e. two negotiating agents), and clauses were limited to binary values, meaning that the clause was or not present in a given contract. Even with such restrictions, the domain of the solution space may become very large. For instance, a negotiation scenario with 50 possible clauses would yield a search space of about 10^{15} possible contracts. This, along with the assumption of non-linearity in the agents' preference spaces, imposed serious difficulties for the negotiation. First, agents needed to use nonlinear optimization mechanisms to try to find desirable contracts within their own preference spaces. Once desirable contracts for each agent were identified, building agreements had its own difficulties, since the scenario was assumed competitive, and thus agents were not inclined to fully disclose their preferences.

Though there are negotiation scenarios about complex contracts which may be modeled with such a solution space, in many cases more than two agents are involved in a negotiation. Also, most contracts may have non-binary clauses. In a rental agreement, for instance, clauses may state the rent, the security deposit or the length of the lease. A labor agreement may include different insurance options. Such issues may have a larger domain, which can greatly increase the solution and preference space complexity.

Taking this into account, in this work we focus in the general case of multilateral negotiations of complex contracts, where the issues or clauses included in the contracts have discrete domains. We also assume that agents' preferences about the different issues are not independent, which means that the utility that a given clause in the contract yields for an agent may depend on the presence of other clauses. Interdependence between attributes in agent preferences can be described by using different categories of functions, like K-additive

Fig. 2 Example of a utility space with two issues and three constraints



utility functions [8,22], bidding languages [61], or weighted constraints [31]. In this work, we model this dependency in agent preferences by means of weighted constraints, which are a natural way to model user preferences and to express dependencies between issues [51]. These constraints may represent ranges of values over different issues, meaning that when *all* the clauses affected by the constraint have values satisfying it, that yields a given utility value for the agent. The set of an agent’s constraints and their associated utility values builds the preference space of the agent.

From a geometrical point of view, 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 than or equal to 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 2 shows an 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 $C2$ (that is, constraints $C1$ and $C2$ overlap, creating a region of higher utility). Contract y , on the other hand, 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 negotiation domain and an agent’s preference model by means of a set of definitions:

Definition 3 *Issues under negotiation.* The *issues under negotiation* are defined as a finite set of variables $X = \{x_i | i = 1, \dots, n\}$.

Definition 4 *Solution space.* The negotiation *solution space* is defined by the values that the different values may take. To simplify, we assume that issues take values from the domain of integers $[0, x_D^{\max}]$:

$$D = [0, x_D^{\max}]^n$$

Definition 5 *Contract or potential solution.* A contract or potential solution to the negotiation problem is a vector $s = \{s_i | i = 1, \dots, n\}$ such that $s \in D$ defined by the issues' values.

Definition 6 *Constraint.* A constraint is a set of intervals which define the region where a contract must be contained to satisfy the constraint. Formally, a constraint c is defined as

$$c = \{I_i^c | i = 1, \dots, n\},$$

where $I_i^c = [x_i^{\min}, x_i^{\max}]$, with $x_i^{\min}, x_i^{\max} \in [0, x_D^{\max}]$ 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.

Definition 7 *Constraint satisfaction.* 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 $s \in x(c)$, meaning that s is in the set of contracts that satisfy c .

Definition 8 *Preference space.* An agent's preference space may be defined as a tuple

$$\langle C, \Omega \rangle,$$

where $C = \{c_k | k = 1, \dots, l\}$ is a set of constraints over the values of the issues x_i for the agent and $\Omega = \{\omega(c_k) | k = 1, \dots, l; \omega(c_k) \in \mathbb{N}^+\}$ is a set of weights or utility values, such that $\omega(c_k)$ is the associated utility value for constraint c_k . For simplicity, we will assume that constraint weights take values from the set of positive integers.

Definition 9 *Utility function.* An agent's utility function 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 [31]. As we have seen in Sect. 2.2, the degree of complexity of the utility spaces produced depends on the number of issues, the domain of the issues and the structural properties of the utility spaces. For the purpose of this work, we make the following assumptions:

- We assume that the number of issues and the domains of the issues are such that they make exhaustive search within the utility space of the agents intractable.
- We assume that the utility spaces of the agent are highly uncorrelated, and so no *a priori* assumptions may be made about where high utility contracts may be located. Therefore, agents may need to resort to local nonlinear optimization techniques to identify such high-utility contracts.
- We assume knowledge about other agent's preferences not to be common (i.e. agents do not know their opponent preference structures, neither they can compute opponent's utility for a given contract).
- We assume that the negotiation setting is competitive, and that agents may be unwilling to reveal too much information about their preferences to the other negotiating agents.

The negotiation protocol and mechanisms proposed, which are described in the next sections, are specifically designed to address this negotiation setting. However, through the study performed in the latter sections of this paper, some of the assumptions are relaxed to evaluate the influence of agent strategies and variations in the correlation lengths of the utility spaces over the negotiation outcomes.

3.2 Interaction protocol

As we stated at the beginning of this section, our model relies on a mediated, auction-based protocol to support agent interaction. The reason for the choice of such a protocol is two-fold. On one hand, the auction-based approach allows to efficiently cope with many of the challenges imposed by multilateral interactions [39]. On the other hand, the use of a mediator allows to decouple individual agent goals (maximizing their own payoff) from social negotiation goals (usually, reaching an agreement which maximizes social welfare). This makes easier mechanism and strategy definition, since agents can be assumed selfish and competitive, while the mediator can be entitled with the more-cooperative task of pursuing social welfare.

Since the main focus of this work is on agent strategic behavior, we have chosen a simple, one shot, auction based interaction protocol for the negotiation, which mainly consists of two steps:

1. *Bidding*: Each agent j generates a set of n_b^j bids $B^j = \{b_i^j | i = 1, \dots, n_b^j\}$, where each bid b_i^j represents a region within the solution space which only contains contracts that agent j would be willing to accept as solutions. Each agent sends its bid set B^j to the mediator, along with the utility associated to each bid.
2. *Deal identification*: The mediator tries 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. A final deal is chosen from the set of potential solutions, according to social welfare criteria.

The protocol, as described, is fairly straightforward, and the decision mechanisms which agents employ for bidding and deal identification are the ones which mostly determine the effect of agent strategic behavior. There are many different mechanisms which can be used in this context. In the following we briefly describe the ones we have found to yield better results in terms of negotiation efficiency and failure rate. All these mechanisms rely on the concept of *quality factor*, which we introduce in the following section.

3.3 Constraint/bid quality factor

The use of weighted constraints generates a “bumpy” utility space, with many peaks and valleys. However, the degree of “bumpiness” is highly dependent on the way the constraint set is generated, and specially on the average width of the constraints. Figure 3 shows an example of the resulting two-dimensional utility space for 50 binary constraints, where the domain of the issues is chosen to be $[0,9]$, and constraints are generated by choosing the width of each constraint in each issue randomly within the $[3,7]$ interval. This generates rather “wide” constraints. On the other hand, Fig. 4 shows an utility space obtained using “narrow” constraints, choosing their widths from the $[1,2]$ interval. Comparing both figures we can see that, though both utility spaces are nonlinear, the space generated using narrow constraints is more complex, with narrower peaks and valleys. As the number of issues under consideration increases, the differences between having wide or narrow constraints become more relevant. For instance, the average correlation length for utility spaces generated using $[3,7]$ constraints for six issues is $\psi = 5.9$, while average correlation length for utility spaces generated using $[1,2]$ constraints is $\psi = 2.8$. Though most utility-maximizing negotiation approaches work in scenarios like the example shown in Fig. 3, their performance (in terms of optimality and failure rate) decreases drastically in highly nonlinear scenarios defined using

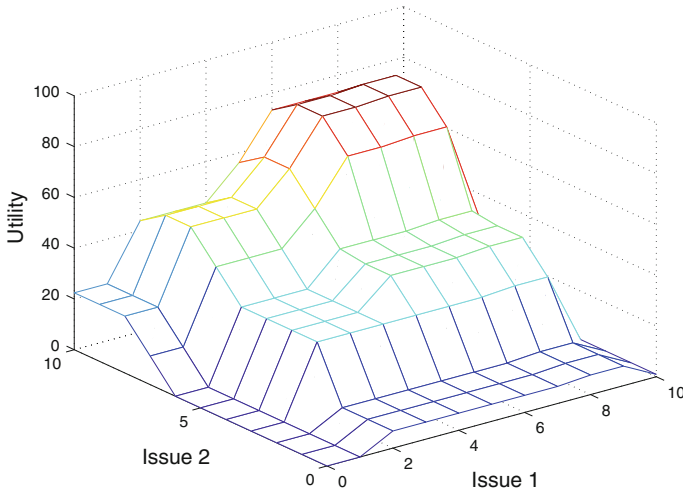


Fig. 3 Example of a nonlinear utility space generated by using “wide” constraints




Fig. 4 Example of a highly uncorrelated utility space generated by using “narrow” constraints

narrow constraints, and therefore an alternative approach is needed to deal with these highly uncorrelated utility spaces [55].

If we compare the utility spaces shown in Figs. 3 and 4, we can see that the main difference between them (apart from the absolute utility values, but they have no effect in optimality) is the width of the peaks. Highly-nonlinear scenarios will yield narrower peaks. Utility maximizing agents tend to choose those peaks (or high-utility regions) as bids, and 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), 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. Intuitively, in such complex scenarios, an agent with no knowledge of the other agents’ preferences should deviate from the “plain utility maximization strategy” and 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).

We formally represent this through the following definitions:

Definition 10 *Volume of a region.* The *volume* of a given region r within the solution space D (be it a constraint or a bid) is defined as the cardinality of the set of contracts contained within the region.

$$v_r = |r|, \quad \text{with } r \subset D$$

Definition 11 *Quality factor.* The *quality factor* of a given region r within the solution space D (be it a constraint or a bid) is defined as

$$Q_r = u_r^\alpha \cdot v_r^{1-\alpha},$$

where u_r and v_r are, respectively, the utility and volume of the bid or constraint r , and $\alpha \in [0, 1]$ is a parameter which models the attitude of the agent. A social, cooperative or 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, highly competitive or selfish agent ($\alpha > 0.5$) will, in contrast, give more importance to bid utility.

3.4 Bid generation mechanisms

3.4.1 Contracts sampling and simulated annealing

We can see the problem of finding the adequate set of bids for an agent as a local optimization problem, since for rational agents bids should be high-utility regions or, more generally, regions of high quality factor. Therefore, nonlinear optimization mechanisms may be used by the agents to find those regions suitable to be sent to the mediator as bids. Here we describe a bidding mechanism based on simulated annealing, which consists of three 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. The function which is tried to maximize by the simulated annealing optimizer is the quality factor Q . Since the quality factor Q is a feature of a region, not a contract, the adjusted contracts must be mapped to the high utility regions where they are contained before they are accepted or rejected by the simulated annealing engine. This can be easily done by checking all constraints in the agent preference model and computing the intersection of the constraints which are satisfied by the candidate contract. The volume of this intersection can then be used to compute the quality factor Q of the region. This results in a set of high-quality contracts.
3. *Bidding:* Each agent generates a bid for each high-quality, adjusted contract. The bids are generated as the intersection of all constraints which are satisfied by the contract. Bids defined in this way represent hyper-rectangle regions in the n -dimensional solution space. Each agent sends its bids to the mediator, along with the utility associated to each bid.

The bid generation mechanism may be seen formally in Algorithm 1. Also, some details about the mechanism are highlighted. The algorithm is run for a fixed number of iterations n_b , which imposes the maximum number of generated bids (1). The function *adjust_annealing* ($x, Q(\cdot, \alpha), n_{SA}, T_{SA}$) uses simulated annealing to return a region of optimal quality factor

using as starting point a sampled contract x (2). There are some parameters in this function which may be adjusted to influence the behavior of the simulated annealing algorithm, like the initial temperature and the number of iterations. As studied in [30], best results in term of optimality and efficiency are achieved using $n_{SA} = 30$; $T_{SA} = 30$. Moreover, the algorithm discards any contract which, once adjusted, yields less utility than the agent's reservation value u_R , which guarantees that all bids would be accepted by the agent as final solutions (3). Finally, duplicate bids or bids contained in other bids are also discarded (4). Some of these ideas are also used in the next bidding mechanism described.

Algorithm 1: Bid generation using simulated annealing over quality factor

Input:
 D : solution space domain
 n_b : maximum number of bids
 u_R : reservation utility for the agent
 C : constraint set defining agent's utility space
 u : agent's utility function
 α : agent's attitude parameter
 Q : function which computes the quality factor of a region
 n_{SA} : iteration bound for the simulated annealing algorithm
 T_{SA} : initial temperature for the simulated annealing algorithm

Output:
 B : bid set
 $B = \emptyset$;
 $k = 0$;

```

1 while  $k < n_b$  do
  |  $k = k + 1$ ;  

  |  $x = \text{random\_contract}()$ ;  

2    $b = \text{adjust\_annealing}(x, Q(\cdot, \alpha), n_{SA}, T_{SA})$ ;  

3   if  $u(b) \geq u_R$  then  

  |    $B = B \cup b$ ;  

  |  

  | end  


4 remove_duplicates( $B$ )
```

3.4.2 Maximum weight independent set and the max-product algorithm

There have been a number of recent successful efforts in literature for using graphs to model negotiation scenarios in multi-link negotiations [89] or combinatorial auctions [21]. One of the advantages of such approaches is that they allow to use well-known graph methods for solving the negotiation problem. In our case, graphs provide 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 Fig. 2. 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 Fig. 5.

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

Fig. 5 Weighted undirected graph resulting from the utility space in Fig. 2



(MWIS). Though MWIS problems are NP-hard, in [3], a message passing algorithm is used to estimate MWIS. The algorithm is a reformulation of the classical max-product algorithm called “min-sum”, and works as follows. Initially, every nodes i send their weights ω_i to their neighbors $N(i)$ as messages. At each iteration, each node i updates the message to send to each neighbor j by subtracting from its weight ω_i the sum of the messages received from all other neighbors except j . If the result is negative, a zero value is sent as message. Upon receiving the messages, a node is included in the estimation of the MWIS if and only if its weight is greater than the sum of all messages received from its neighbors. Message passing continues until MWIS converges or the maximum number of iterations is exceeded. This is formally shown in Algorithm 2.

Algorithm 2: Min-sum algorithm for MWIS estimation

Input: $i = 1, \dots, n$: nodes (constraints) in the weighted graph $\omega_i | i = 1, \dots, n$: weight (utility) of each node (constraint) $N(i)$: set of neighbors of each node (incompatible constraints) t_{\max} : maximum number of iterations

Output: MWIS: estimation of the MWIS

$t = 0; m_{i \rightarrow j}^t = \omega_i \forall j \in N(i)$ **while** $t < t_{\max}$ **do**

$t = t + 1$; **foreach** i **do**

$m_{i \rightarrow j}^t = \max\{0, \omega_i - \sum_{k \neq j, k \in N(i)} m_{k \rightarrow i}^{t-1}\}$

end

$MWIS^t = \{i | \omega_i > \sum_{k \in N(i)} m_{k \rightarrow i}^{t-1}\}$ **if** $t > 1$ **and** $MWIS^t = MWIS^{t-1}$ **then**

return $MWIS^t$

end

However, this reformulation of the bidding problem is not in itself a suitable solution, since it has some serious drawbacks. On one hand, the algorithm is deterministic, and thus only one bid can be generated for a given set of constraints. On the other hand, the algorithm is based on utility maximization, so it does not allow the agent to search for high quality bids. Moreover, the quality factor Q cannot be directly introduced into the max-product or min-sum algorithm, because the algorithm is based in a weighted graph where weights are additive, and the quality factor is not additive (that is, the quality factor of the intersection of a set of constraints is not the sum of the quality factor of the constraints).

To solve this, the algorithm is applied to a subset of constraints $C' = \{c'_k | k = 1, \dots, n_c; n_c < l; c'_k \in C\}$. 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. The approach proposed in can be seen in Algorithm 3. In order to maximize quality factor of the generated bids, a *tournament selection* [57] is used when

generating the subset of constraints C' to be passed to the max-product algorithm (1). This tournament selection works as follows. For each bid to generate, a number n_t of candidate constraint subsets are randomly generated. From these subsets, the one which maximizes the product of the quality factors Q of its constraints is chosen as the subset C' to be used for the max-product algorithm. In this way, since high- Q constraints are more likely to be selected, we expect the average Q for the resulting bids to be higher.

Algorithm 3: Bid generation using MWIS and Q -based tournament selection

Input:

n_b : maximum number of bids
 u_R : reservation utility for the agent
 C : constraint set defining agent's utility space
 Ω : constraint weights for the agents
 u : agent's utility function
 n_c : number of randomly chosen constraints passed to the MWIS algorithm
 n_{MWIS} : maximum number of iterations for the MWIS algorithm
 α : agent's attitude parameter
 n_t : number of candidate subsets in tournament selection

Output:

B : bid set

$B = \emptyset$;

$k = 0$;

while $k < n_b$ **do**

$k = k + 1$;

1 $C' = \text{tournament_selection}(C, n_c, \alpha, n_t)$;

$\{\text{nodes, weights, neighbors}\} = \text{build_tree}(C', \Omega)$;

$MWIS = \text{minsum}(\text{nodes, weights, neighbors, } n_{MWIS})$;

$b = \text{generate_bid}(C', MWIS)$; **if** $u(b) \geq u_R$ **then**

$B = B \cup b$;

end

$\text{remove_duplicates}(B)$

3.5 A probabilistic mechanism for deal-identification

Once agents have placed their bids, it is the turn to the mediator to try to find deals among them. The most straightforward way to do this is to perform an exhaustive search of overlaps between the different agents' bids, tagging those overlaps found as potential solutions, and then selecting a winner solution from the potential solution set according to social welfare criteria.

The problem with such an exhaustive search is scalability with the number of agents. In a worst case scenario, the mediator would have to search through a total of $n_b^{n_a}$ bid combinations, where n_b is the number of bids per agent, and n_a is the number of negotiating agents. This imposes a limit on the maximum number of bids that an agent may send to the mediator. For instance, if we limited the number of combinations to 6,400,000, this means that, for four negotiating agents, the maximum number of bids per agent is $\sqrt[4]{6400000} = 50$. This limit becomes harder as the number of agents increases. For example, for ten agents, the limit is four bids per agent, which drastically reduces the probability of reaching a deal. This is specially true for highly-nonlinear utility spaces, where the bids are narrower.

To address this scalability limitation, we perform a probabilistic search in the mediator instead of an exhaustive search. This means that the mediator will try a certain number n_{bc} of randomly chosen bid combinations, where $n_{bc} < n_b^{n_a}$. In this way, n_{bc} acts as a performance parameter in the mediator, which limits the computational cost of the deal identification phase. Of course, restricting the search for solutions to a limited number of combinations may cause the mediator to miss good deals. Taking this into account, the random selection of combinations is biased to maximize the probability of finding a good deal. Again, the parameter used to bias the random selection is Q , so that higher- Q bids have more probability of being selected for bid combinations at the mediator.

The mechanism is formally shown in Algorithm 4. We can see that the number of analyzed bid combinations is limited to n_{bc} (1), and that the function *combine_bids* (...) selects the bid combinations to analyze (2). Limiting bid combinations at the mediator allows us to remove the limit on the bids issued by the agents, which increases the probability of finding potential deals. Finally, the algorithm selects from all deals found the one which maximizes social welfare, computed using the *sw* (s, U) function (3). Social welfare is computed as the Nash product [60], that is, the product of the utilities that a potential solution gives to every agent.

Algorithm 4: Probabilistic deal identification

Input:

A : set of negotiating agents
 $n_a = |A|$: number of negotiating agents
 B : set of bids issued by every agent
 U : declared utilities for every agent's bids
 Q : declared quality factors for every agent's bids
 sw : social welfare function
 n_{bc} : maximum number of bid combinations at the mediator

Output:

s_f : final deal
 $n = 0$;
 $S = \emptyset$;
1 while $n < n_{bc}$ **do**
2 $s = \text{combine_bids}(A, n_a, B, U, Q)$;
 if $s \neq \emptyset$ **then**
 $S = S \cup s$;
 $n = n + 1$;
end
3 $s_f = \arg \{\max_{s \in S} sw(s, U)\}$;

3.6 Discussion

We approach the negotiation problem as a mechanism design problem, where we aim to design the structure of the game in a way that facilitates social welfare optimizing outcomes [58]. We assume a complex agent preference space, where exhaustive search for high-value solutions is unfeasible for the agents. Therefore, preference revelation is performed in the form of bids, which are subsets of the preference space. In fact, the bidding process is seen as a local constraint-based optimization problem, where each agent needs to find combinations of compatible constraints which maximize its own utility. Analogously, the deal identification process is seen as a constraint-based multi-objective optimization problem, where the

mediator tries to find overlaps between agents' bids which maximize social welfare. We have chosen a mediated approach for the negotiation to facilitate social-welfare maximizing mechanism design, and we have used a heuristic search at the mediator to cope with the scalability problems imposed by the high cardinality of the solutions space.

The experimental evaluation performed in our previous work showed that the use of the quality factor in the bidding and identification mechanisms described significantly improved the performance of the negotiations over the previous approaches in highly uncorrelated utility spaces [55]. Furthermore, it was also pointed out that the use of the quality factor greatly improved the scalability of the model, allowing to perform negotiations with up to 14 agents and 20 issues while keeping high optimality values and low failure rates. However, there are some issues which are not addressed in the work. Even when the quality factor is designed to model the attitude of an agent (be it risk attitude, tendency to cooperation or selfishness) through its α parameter, the experimental evaluation was performed only for $\alpha = 0.5$. This assumes that all negotiating agents have the same attitude, and also that this 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. The parameter α allows an agent to take a given strategy (a given attitude), and so the possibility arises that different agents may choose different strategies for a given negotiation.

Since our aim is to design mechanisms which facilitate social welfare optimizing outcomes, we have to pay attention to the consequences of having agents playing different strategies in the negotiation. It could be the case that the proposed approach favored a specific strategy (or set of strategies) against the others. Assuming the agents are individually rational, they would have the incentive to play these favored strategies. If they are different from the assumption above, the outcomes of a real negotiation among rational agents could differ from the ones obtained in our previous experiments in terms of social welfare. Therefore, a strategy analysis is needed to evaluate the mechanisms in situations where agents with different attitudes interact.

4 Strategy analysis of the auction-based negotiation protocol

As we stated in Sect. 2.1.3, one of the main challenges when designing decision mechanisms for automated negotiations is strategic stability, and this problem is closely related to the notions of *equilibrium* described above. For heuristic approaches such as those described above, game theory concepts and analyses cannot be directly applied, due to the high variability of the bid generation mechanisms and the total uncertainty about the preferences of the different agents. There are some successful works for finding equilibrium conditions under incomplete information [24, 81], and even with infinite games [68]. However, all these works assume a certain degree of determination about the outcome of the negotiation once the agents (each one having a private type) have chosen their strategies. With pure strategies, this determination is perfect, that is, negotiation outcome is known as soon as agents have chosen their strategies. For mixed strategies, agents have a probability distribution over their set of possible actions, and thus the outcome of the negotiation is not perfectly determined until all agents have chosen their actions.

In the heuristic approach we are dealing with, there are many levels of uncertainty. Agent strategies may be modeled by varying the value of the α parameter used to compute quality factor. This can be seen as a pure strategy, since the choice of an agent is to use one value of α or another. However, a negotiating agent final action (i.e. the bids which are actually sent to the mediator) does not depend only on that choice. It also depends, of course, on the

agent’s preference model, which may be identified with the agent “type”. However, since agents do not know or fully explore their utility spaces (we assume that such exploration is computationally intractable), the final agent action also depends on the heuristic search method used to generate the bids. Since this method is, in the cases outlined in the previous section, non-deterministic, this adds an additional layer of uncertainty, which we could in some way identify with the use of mixed strategies (although very complex ones). In addition, once all negotiating agents have performed their actions (i.e. bids), the mediator initiates the deal identification step of the protocol, which is also non-deterministic. These multiple layers of uncertainty make very difficult to directly apply game-theoretic concepts such as equilibrium conditions or best-response strategies, since different trials of the same “game” (same agents, same strategy combinations, same preference sets) may yield drastically different results depending on the specific outcomes of the heuristics involved. Therefore, part of our study would be necessarily empirical, which is a usual approach when dealing when heuristic strategies [1].

Some of the game theory concepts, however, can still be useful with some nuances. In particular, strategic properties analogous to the equilibrium conditions in game theory may be studied for heuristic mechanisms. This section is dedicated to assess the strategic behavior of the described auction-based negotiation model, determining the existence of individually optimal strategies and social optimal strategies, and verifying if the auction-based negotiation mechanisms are prone to situations involving high values for the price of anarchy (PoA). To this end, a probabilistic analysis and an empirical evaluation have been performed.

4.1 Probabilistic analysis

Intuitively, it can be seen that the quality factor defined above 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 can be found in Appendix A. For the purpose of this section, the final expressions will suffice. In particular, deal probability for a single run of the auction-based negotiation protocol is given by

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

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

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

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 [55]. Of course, this strategy does not model the attitude of, for instance, 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 or a selfish agent could 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 attitude introduced in the quality factor in the previous sections. 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, or it is known that the utility spaces of the different agents are strongly different, 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, or be more cooperative, or less selfish) 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 selfish strategy if the other agents are more cooperative (their lower α values would compensate the decrement in deal probability). However, should all agents decide to use selfish strategies, deal probability would reduce drastically, leading to low expected individual and social welfares. If there is a tendency or incentive for this condition to occur, we would have a high price of anarchy situation, and we should design and establish mechanisms to stabilize the protocol.

4.2 Experimental analysis

In this section the strategic properties of the protocol inferred from the statistical analysis are empirically verified. To this end, a set of experiments has been devised to analyze the main strategic properties of the model. As stated in Sect. 2.1.3, these properties are related to the different notions of equilibrium. However, as we discussed above, determining rigorous equilibrium conditions in our negotiation model is very difficult, due to the different layers of uncertainty introduced by the heuristics used. Therefore, the experiments performed and the conclusions drawn from them will be based on statistical observations, in a similar way to the notions of equilibrium considered for Bayesian players in Harsanyi [24] and Reeves and Wellman [68]. In particular, best-response strategies will be determined according to the maximization of the expected payoff.

To conduct the experiments, negotiating agents will generate their offers using contract sampling with Q-based simulated annealing (SA-Q) or maximum weight independent sets with a Q-based tournament selection (MWIS-Q). The experiments have been designed to study the dynamics of the negotiation process when agents with different strategies interact. In this context, agent strategic behavior is defined by the value of the α parameter each agent

uses to compute constraint and bid quality factor. Preliminary versions of some of the results included in this section have been previously published in Marsa-Maestre et al. [56].

4.2.1 Individually optimal strategy analysis

First of all, the existence of an *individually optimal strategy* is studied. This is closely related to the concept of *dominant strategy* defined in game theory. A dominant strategy would be one which, regardless of the strategies the other agents choose, ensures that a given agent would not have achieved a higher payoff using any other possible strategy. However, in a model with such degree of variability in bid generation and deal identification as the one we are dealing with, and with infinite strategies for the different agents (the possible values for the α parameter), it is not possible to achieve this certainty. In particular, it is not possible to state that a given strategy *would have given* an agent a better payoff than another strategies, since the same strategies may yield drastically different payoffs in different trials. We can, however, evaluate statistically which strategies tend to give the agents the best payoffs, trying to determine whether there is a tendency in the model to favor a given strategy. This would be an individually optimal strategy in the context of our heuristic model.

Though the idea of an individually optimal strategy is conceptually simple, evaluating its existence is not straightforward. At a first glance, we need to be able to compare the utilities or payoffs obtained by an agent in different trials of the experiment. However, to see if there is an individually optimal strategy regardless of the agent's specific preferences, payoffs obtained by agents with different preference spaces need to be evaluated too. The problem is not only that maximum potential payoffs for different agents may vary, but also that such potential payoffs for a given negotiation encounter also depends on the preference spaces of the *other* agents participation in the negotiation, since only those regions of the solution space whose utility is above the reservation values of *all* agents are actual potential solutions. Taking this into account, we measure the payoff obtained by a given agent j on a given encounter as its *individual optimality rate* defined as the ratio between the payoff obtained by the agent in the encounter, and the highest possible payoff for that agent in that encounter. This highest possible payoff is computed by giving all information about the agents preferences to a nonlinear optimizer, which then computes an approximate optimal contract for j with complete information.

In a first set of experiments, we have tried to determine if there is a strategy, determined by a certain α value, which yields maximum utility to an agent given the strategies of the other agents. To evaluate this, we have performed a set of experiments comparing the utility obtained by an *individualist agent*, which plays an individual strategy determined by α_i , with the utility obtained by the other agents. To model the joint effect of the behavior of the rest of the agents, we have used a common strategy α_s for them. Experiments have been performed varying α_i and α_s within the interval $[0, 1]$ in 0.1 steps.

Figures 6 a and b show the box plots of the results for 100 runs of the experiments for SA-Q and MWIS-Q, respectively, for six agents and six issues. We have represented the ratio between the optimality rates obtained by the individualist agent and the utility obtained by the rest of the agents. In this case we consider only successful negotiations, since in failed negotiations all agents get zero utilities, and the ratio cannot be computed. We can see the same trend for both approaches studied. Generally, the individualist agent obtains a higher utility when using higher α_i values. We can also see that, for any α_s , the maximum utility value for the individualist agent is obtained for $\alpha_i = 1$, which suggests that this could be the individually optimal strategy. For $\alpha_s > 0.8$ negotiations failed, and thus no values are

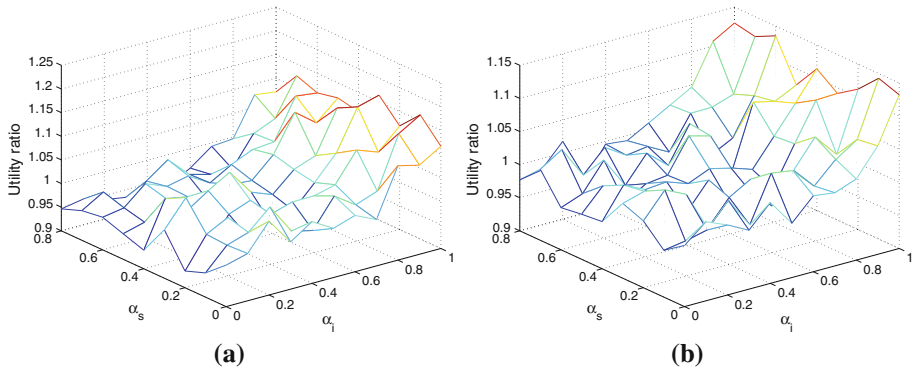


Fig. 6 Individual optimal strategy analysis against symmetric strategy combinations. **a** SA-Q, **b** MWIS-Q

shown in the figures. This result is directly related to social strategy analysis, and thus we will discuss it in more detail in the following section.

Though the results suggest that $\alpha_i = 1$ is the individually optimal strategy for the agents, the previous experiment only tests agent individual strategies against symmetric strategy combinations (i.e. all the other agents play the same strategy). In a more realistic setting, we may expect agents to play non-symmetric strategy combinations. To determine the expected payoffs of the different individual strategies for the individualist agent against arbitrary strategic combinations of its opponents, we have repeated the previous experiment randomizing the strategy choice of the other agents. In this way, the individualist agent played its individual strategy α_i , while the other agents' strategies were randomly drawn from a discrete uniform distribution within the interval $[0, 1]$ in 0.1 steps. Since the use of non-symmetric strategy profiles for the opponents increased the variability of the experiment, 1000 runs of each experiment were performed.

Figures 7 a and b show the box plots of the results for SA-Q and MWIS-Q, respectively, for six agents and six issues. We have again represented the ratio between the optimality rates obtained by the individualist agent and the utility obtained by the rest of the agents. In this case, the columns in the horizontal axis represent the different values for $\alpha_i = 1$, while in the vertical axis we have represented the ratio between optimality rates as notched box and whisker plots. The box and whisker plots are represented as follows. Each column corresponds to a set of samples of the gain for individualist agents in 100 negotiations. The two boxes in each column contain 50% of the samples, corresponding to the 25th and 75th percentiles, and the red line in the separation of the two boxes represents the median. The small notches around the median display the variability of the median between samples as 95% confidence intervals, computed using the method described in [85]. This means that two medians are significantly different at the 5% significance level if their notches do not overlap. The whiskers (dashed lines) extend to the most extreme data points not considered outliers, and outliers are plotted individually with a plus (+) sign. We can observe similar results than in the previous experiment. The individualist agent obtains a higher expected relative payoff when using higher α_i values, being $\alpha_i = 1$ the strategy maximizing expected payoff, so we can conclude that this is the *individually optimal strategy* for the agents.

4.2.2 Social strategy analysis

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

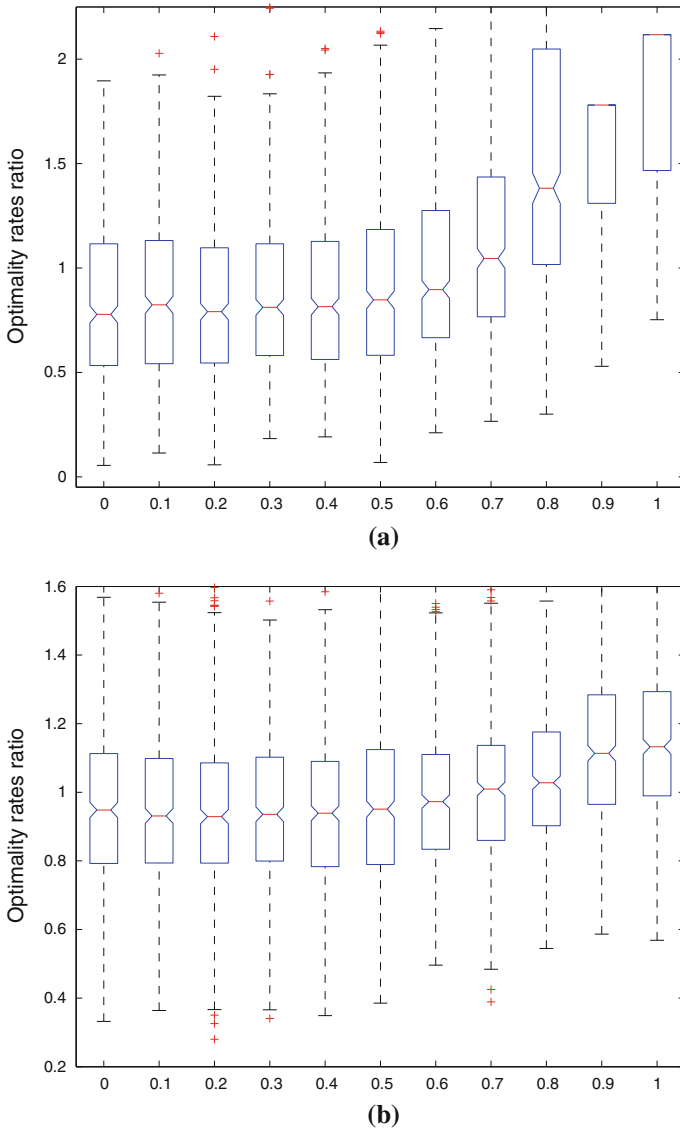


Fig. 7 Individual optimal strategy analysis against random strategy combinations. **a.** SA-Q. **b.** MWIS-Q

expected 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 varying 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 correlation length ψ , as introduced in Sect. 2.2.

Table 1 Social strategy analysis for SA-Q

ψ	α_s										
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
2.8	0.327	0	0	0	0	0	0	0	0	0	0
3.1	0.529	0	0	0	0	0	0	0	0	0	0
4.0	0.772	0	0	0	0	0	0	0	0	0	0
4.3	0.864	0.884	0.897	0.830	0.867	0.907	0.919	0.935	0.948	0	0
4.6	0.935	0.955	0.959	0.961	0.963	1.000	1.000	1.000	1.000	1.000	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

Table 2 Social strategy analysis for MWIS-Q

ψ	α_s										
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
2.8	0.334	0.379	0.384	0.377	0.434	0.480	0.552	0.486	0	0	0
3.1	0.460	0.528	0.495	0.504	0.554	0.555	0.596	0.682	0	0	0
4.0	0.795	0.785	0.798	0.814	0.821	0.838	0.828	0.827	0.814	0	0
4.3	0.967	0.963	0.976	0.961	0.973	0.969	0.971	0.970	0.977	0	0
4.6	1.000	1.000	0.975	1.000	1.000	1.000	1.000	1.000	1.000	1.000	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

Experiment results for six agents and six issues for SA-Q and MWIS-Q are presented, respectively, in Tables 1 and 2. Each table shows the median *social optimality rates* for the negotiation as the value of α_s varies, for different values of ψ . Social optimality rate is defined as the ratio between the social welfare obtained with the protocol and the social welfare obtained using an optimizer with complete information. For SA-Q, in the most uncorrelated utility spaces, only the most risk-averse strategy ($\alpha_s = 0$) achieves successful negotiations. For medium or low-complexity scenarios, maximum social welfare values are obtained for α_s values around 0.7. MWIS-Q approach performs better than SA-Q for uncorrelated utility spaces, and 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 optimal social strategies have been identified, a desirable property would be that these strategies were a *Nash equilibrium* or a *Bayes-Nash equilibrium* for the system as we saw in Sect. 2.1.3, that is, that there was no incentive (no potential increase in expected payoff) for any agent to deviate from this strategy. Unfortunately, as we saw above, there is an individually optimal strategy, given by $\alpha_i = 1$. Therefore, an individually rational agent may decide to take this strategy to maximize its own benefit (as seen in Fig. 6 a and b). All agents have the same incentive, so the trend would be for all agents to choose $\alpha_i = 1$. As we can see in Tables 1 and 2, this makes negotiations fail in medium and highly complex scenarios. The fact that individual rationality may lead the system to situations far from the social optimum makes the model prone to situations analogous to those of high price of anarchy (PoA) described in Sect. 2.1.3. Rigorously speaking, we cannot use Price of Anarchy directly,

since it is related to the notion of Nash Equilibrium, which has no sense in our setting due to the great variability and uncertainty about negotiation outcomes. However, other authors have recently defined analogous concepts to PoA for games under uncertainty conditions, like Bayes-Nash PoA in Leme and Tardos [46]. We can take a similar approach under the assumption that agent types are not known and there are no specific *a priori* beliefs about the strategies played by other agents, which means that from the point of view of the agents, opponents' strategies/types are equiprobable. Taking this into account, we analogously define an *Expected Price of Anarchy* as follows:

Definition 12 *Expected Price of Anarchy (EPoA)* The *Expected Price of Anarchy* 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 minimum expected social welfare achievable by means of an *individually rational* agent strategy combination.

$$EPoA = \frac{\max_{s \in S} E[sw(s)]}{\min_{s \in S_{i,r}} E[sw(s)]},$$

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 .

According to this definition and to the results of the experiments above, our negotiation model could be prone to high EPoA situations in medium and highly complex scenarios. If confirmed, this would be a situation which would negative impact model stability. Stability issues in the model, along with techniques to improve stability, are discussed in detail in the following section.

5 Addressing infinite expected price of anarchy 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 address situations of high price of anarchy in the negotiation process are proposed, and their effectiveness is empirically evaluated.

5.1 Enforcing socially-oriented strategies at the mediator

The final element in the deal identification mechanism is the social welfare function $sw(s, U)$. Once a set of viable solutions has been found, the mediator chooses as the solution the one which maximizes social welfare. Therefore, a metric which allows the mediator to compare the different solutions in terms of social welfare is needed. One of the most widely used is usually called *social welfare*, which is defined as the sum of the utilities that solution gives to every agent [67]. Maximizing this metric, solutions near to the Pareto-optimal region are found. However, sometimes the solutions found may have excessive low utility for some of the agents. This is specially true if the agents' reservation value is zero, since there may be solutions maximizing the sum of utilities even when the utility values for some of the agents tend to zero. To avoid this, an alternative metric could be the *minimum utility*, that is, the minimum of the utilities that solution gives to each agent. Though maximizing this metric guarantees a certain satisfaction level for all agents participating in the negotiation, it has an important drawback, since it makes no difference between solutions which give the

same minimum utility even when they give different utility values for the rest of the agents. Therefore, solutions obtained using this criterion may be far apart from the Pareto front.

A metric which allows to achieve more egalitarian solutions which are closer to the Pareto-optimal region is the Nash product [60], which is the product of the utilities that solution gives to every agent. This metric for the quality of a solution is widely used in the literature, since it allows to achieve solutions close to the Nash solution, which is widely used in the literature as a reference for optimality in negotiation processes. The ratio between the Nash product of a given solution to a negotiation problem and the Nash solution associated to that problem is usually referred as the *Nash optimality* of the solution.

Given the different social welfare metrics, it is clear that an agent's attitude greatly influences the final utility value for this agent if an agreement is reached. Once all valid intersections have been found, the final outcome is selected using a function which depends on the utility values the outcome gives to the agents. Selfish, risk-willing or highly competitive agents, which have given more importance to utility against volume in the bid generation process, will have, on average, higher utility bids, and thus their expected deal utility (Eq. 3) will be higher. Taking this into account, the preferred strategy of an agent may be to take a selfish attitude, as we inferred in the previous section. The problem is that, in complex utility spaces, having all agents taking such attitudes could lead to very narrow offers at the mediator, which would make deal probability (given by Eq. 1) decrease drastically. This may lead the protocol to negotiation failures, with zero social welfare, thus resulting in situations of infinite Expected Price of Anarchy, turning the negotiation model unstable.

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 seems reasonable to entitle the mediator with the task of pursuing social welfare. In the deal identification mechanism described in Sect. 3.5, the mediator chooses as the final solution the one maximizing social welfare. The metric used to compute social welfare in this case is 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. This can be done by using a generalized or asymmetrical version of the Nash product, similar to the ones used in [35] to model agents power of commitment. In particular, we propose a modification of the Nash product which we have called *weighted product by average volume*:

Definition 13 *Weighted product by average volume* The *weighted product by average volume* of a solution to a negotiation problem among n_a agents is the product of the utilities the solution gives to every agent i , weighting each utility $u^i(s)$ by an adjustment factor equal to the ratio between the average volume of the bids issued by the agent \bar{v}^i and the maximum average volume of the bids of one of the agents:

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

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, we propose to consider bid utility and bid volume jointly, using a *product weighted by average quality factor*:

Definition 14 *Weighted product by average quality factor* The *weighted product by average quality factor* of a solution to a negotiation problem among n_a agents is the product of the utilities that solution gives to every agent i , weighting each utility $u^i(s)$ by an adjustment factor equal to the ratio between the average quality factor of the bids issued by the agent \bar{Q}^i and the maximum average quality factor of the bids of one of the agents:

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

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

When this last metric is applied, quality factor is not only used to compute social welfare at the mediator. As we have seen in Sect. 3.5, 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 assessment of the bids made by the mediator strongly 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 socially optimal 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 (thus being 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.

Finally, we shall consider that the use of such asymmetrical social welfare metrics, though may contribute to improve model stability, may have its drawbacks as well. The rationale behind the metrics is to “reward” those agents which are playing more cooperative strategies, but the metrics are based on observations about agents' final actions, since their strategies are unknown to the mediator. More specifically, the mediator cannot distinguish whether an agent is issuing low volume or low quality bids because it is playing a selfish strategy or because its utility space does not contain better feasible regions. In this way, the mediator may seem to be giving an undue advantage to agents with wider constraints. This kind of asymmetric models have, however, been used successfully in other negotiation scenarios. The Clarke tax method [11], which was briefly discussed in Sect. 2.1.3 imposes a tax to each agent once the negotiation has ended, making each agent “pay” for the impact that its participation had over other agents' utilities. The approach we have taken here is similar in the sense that we apply the asymmetrical social welfare metrics at the final steps of the deal identification, to select the final deal among all potential deals found, and this final selection

is biased to reward those agents which had issued “better” bids (according to the mediator’s criteria). Though it is not possible to know, with a one-shot negotiation, whether agents are issuing “better” bids due to a socially-oriented strategy or to a more correlated utility space, the expected effect is that agents will have an incentive to select the bids they send not only according to their own utility, but also to the mediator’s criteria, which would be set to favor social welfare.

5.2 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 socially optimal strategy taking a more selfish approach. To evaluate this empirically, we have performed experiments comparing the utility obtained by an *individualist* agent (or, more appropriately, a *selfish* agent, since it seeks to maximize its own payoff), using its individually optimal strategy $\alpha_i = 1$, against the utility obtained by the agent when using the corresponding socially optimal strategy α_s , assuming the rest of the agents are using α_s . Experiments have been made for utility spaces with different correlation lengths. Furthermore, since the model is designed for multi agent negotiations, experiments have been performed for different number of selfish agents, thus studying the effect of possible coalitions or coincidences.

Figures 8 and 9 show the experiment results for SA-Q and MWIS-Q, respectively, for six agents and six issues. Since the aim of the experiment is to study the stability of the proposed protocol, none of the weighted metrics proposed in the previous section has been used, and social welfare is computed at the mediator using Nash product. In addition, the mediator performs deal identification using the bid quality factors declared by the agents (i.e. there is no α_m). The figures show the ratio between the utilities obtained by selfish agents and the utilities obtained when there are no selfish agents for different correlation lengths and different number of selfish agents, for scenarios of different complexity. The horizontal axis represents the number of individualist or selfish agents, while in the vertical axis we have represented the ratio between utilities as notched box and whisker plots. In each figure, column labelled as “0” represents the dispersion of utility gains when there are no selfish agents. We can see that there are only significant gains for selfish agents in medium complexity scenarios. In high-complexity scenarios (Figs. 8a and 9a), the presence of selfish agents makes the negotiations fail, and thus there is no incentive to deviate from the socially optimal strategy. When utility space complexity decreases (Figs. 8b and 9b), we can see that a selfish agent may obtain gains over 40% for SA-Q and 200% for MWIS-Q. Increasing the number of selfish agents makes negotiations fail, thus making unlikely that coalitions will happen. For medium-low complexity scenarios (Figs. 8c and 9c) there is still a significant gain for selfish agents, and this gain increases with the number of selfish agents up to a number of three (coalitions between more agents make negotiations fail). Finally, for the less-complex scenarios (Figs. 8d and 9d), a selfish attitude does not imply a significant gain in utility, since all agents achieve high utility values using the socially optimal strategy. Tables 3 and 4 summarize the results for SA-Q and MWIS-Q, respectively, showing the medians and the 95% confidence intervals for 100 runs of each experiment. 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, because of the existing incentive for agents to deviate from the social optimal strategy to their individually optimal one ($\alpha = 1.0$). As we have seen in Sect. 4.2.2, having all agents deviating to their individually optimal strategy makes the negotiations fail, and thus this situation is the worst

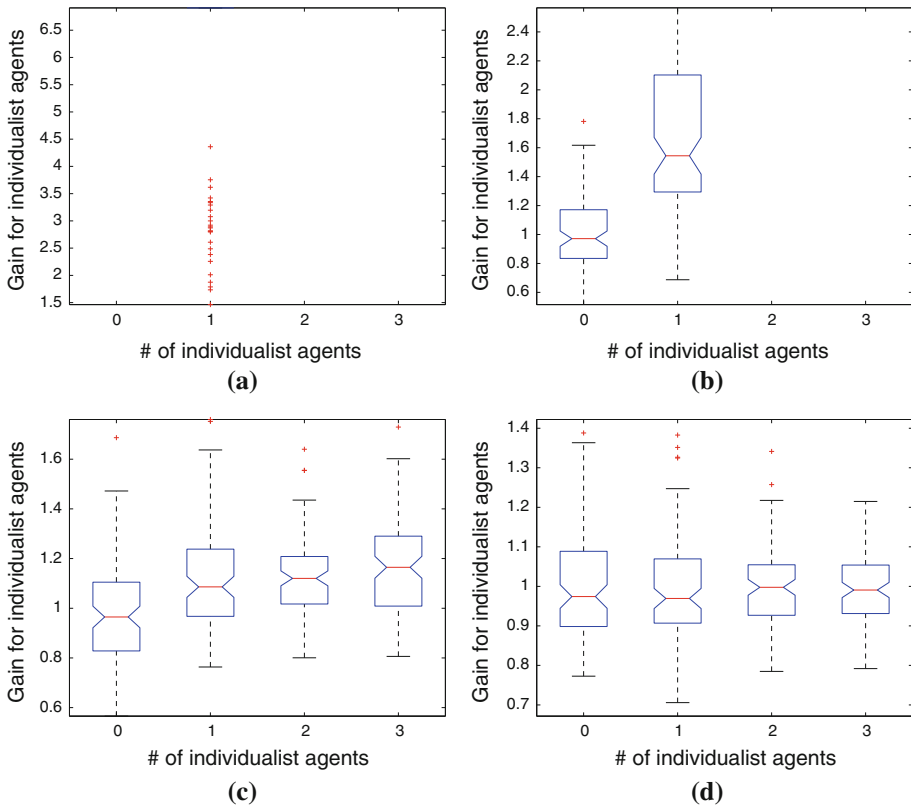


Fig. 8 Stability analysis of the protocol using SA-Q for scenarios with different correlation lengths. **a** $\psi = 2.8$, **b** $\psi = 4$, **c** $\psi = 4.3$, **d** $\psi = 5.9$

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,

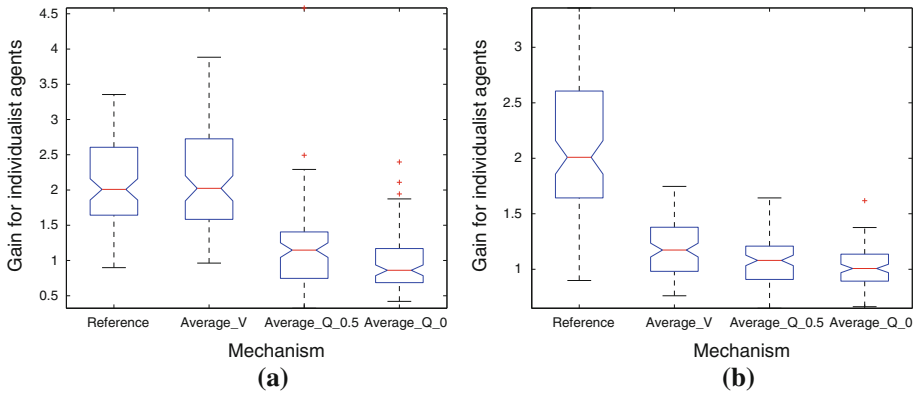


Fig. 11 Effect of the different mechanisms on the stability of the protocol for MWIS-Q in the most critical scenarios. **a** $\psi = 4.0$, **b** $\psi = 4.3$

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 EPoA 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 significative 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

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

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 ascending 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

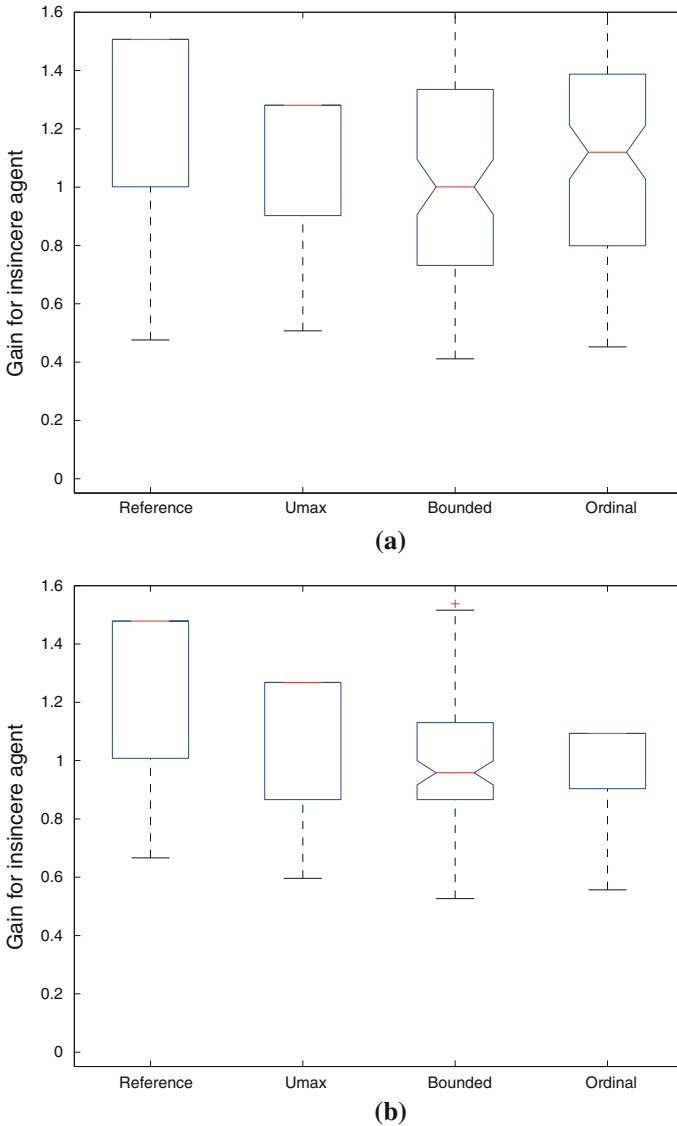


Fig. 13 Effect of the proposed normalization mechanisms for the most critical scenarios. **a** SA-Q, $\psi = 4.0$, 25% degree of exaggeration. **b** MWIS-Q, $\psi = 4.3$, 50% degree of exaggeration

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:

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

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

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 $p(\bigcap_{j=1}^N \overline{E}_j) = (1 - p)^N$. In this case, we can see that the above expression leads to the following:

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

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:

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

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. Chevalyre, 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. *Informati-ca*, 30, 2006.
8. Chevalyre, 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](https://doi.org/10.1007/BF01726210).
12. Ehtamo, H., Hamalainen, R. P., Heiskanen, P., Teich, J., Verkama, M., & Zions, 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](https://doi.org/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/ecolet/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](https://doi.org/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](https://doi.org/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](https://doi.org/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](https://doi.org/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. Sakurai, 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/d228db90c316574e86a4a8ee6278aba>
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-4BN0J81-2/2/4a0255f0b8884b4203ed8c3899fd84e58>
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.