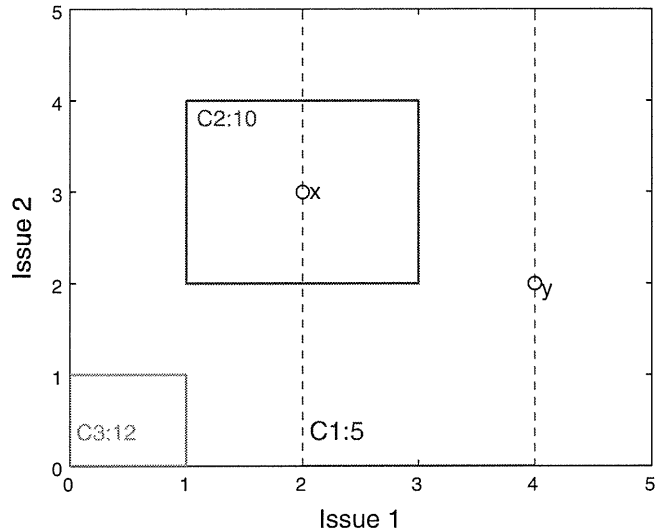


Fig. 2 Example of a utility space with two issues and three constraints



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

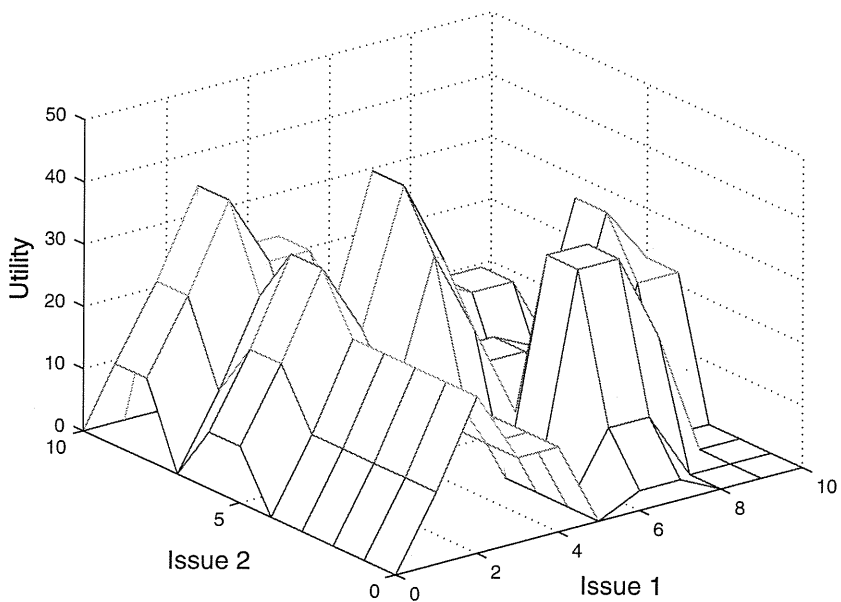


Fig. 4 Example of a highly uncorrelated utility space generated by using “narrow” constraints

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 $\text{remove_duplicates}(B)$

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→j}^t = ω_i ∀ j ∈ N(i) while t < t_max do
  t = t + 1; foreach i do
    | m_{i→j}^t = max{0, ω_i - ∑_{k≠j, k∈N(i)} m_{k→i}^{t-1}}
  end
  MWIS^t = {i | ω_i > ∑_{k∈N(i)} m_{k→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 $C' = \text{tournament_selection}(C, n_c, \alpha, n_t)$; $\{\text{nodes}, \text{weights}, \text{neighbors}\} = \text{build_tree}(C', \Omega)$; $MWIS = \text{minsum}(\text{nodes}, \text{weights}, \text{neighbors}, n_{MWIS})$; $b = \text{generate_bid}(C', MWIS)$; **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 with 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

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

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

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

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

where u_l^k is the utility for the l -th bid of agent k . According to this expression, to maximize expected utility, an agent should reveal as much information as possible. If information disclosure is limited, an agent should try to maximize $\sum_{l=1}^{n_b^k} u_l^k \cdot v_l^k$, balancing in this way bid

utility and bid volume. This is coherent to the choice of $\alpha = 0.5$ in [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

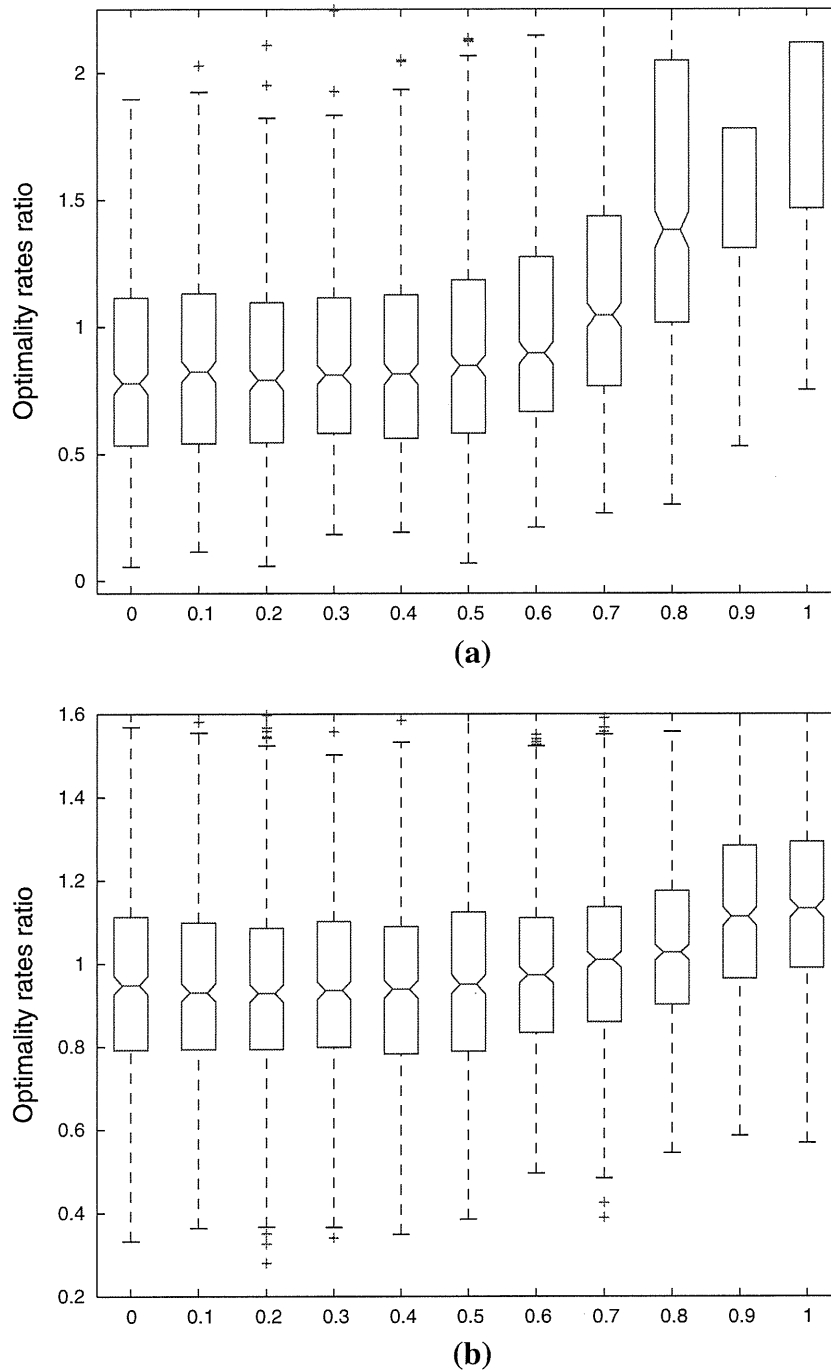


Fig. 7 Individual optimal strategy analysis against random strategy combinations. **a.** SA-Q. **b.** MWIS-Q

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 have a negative impact on 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