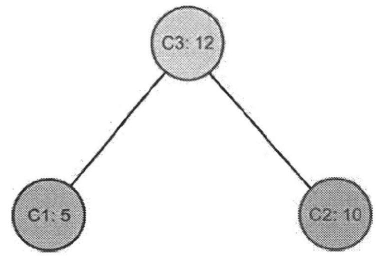

Fig. 5 Weighted undirected graph resulting from the utility space in Fig. 2

(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,...,n: nodes (constraints) in the weighted graph ω_i | i = 1,...,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^t_{i→j} = ω_i ∀j ∈ N(i) while t < t_max do
    t = t + 1; foreach i do
    |   m^t_{i→j} = max{0, ω_i − Σ_{k≠j,k∈N(i)} m^{t−1}_{k→i}}
    end
    MWIS^t = {i | ω_i > Σ_{k∈N(i)} m^{t−1}_{k→i}} 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 $C' = \{c'_k | k = 1, \ldots, n_c; n_c < l; c'_k \in C\}$. 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
$B = \varnothing$;
$k = 0$;
while $k < n_b$ **do**
  $k = k + 1$;
1 $C' = tournament\_selection(C, n_c, \alpha, n_t)$;
  $\{nodes, weights, neighbors\} = build\_tree(C', \Omega)$;
  $MWIS = minsum(nodes, weights, neighbors, n_{MWIS})$;
  $b = generate\_bid(C', MWIS)$; **if** $u(b) \geq u_R$ **then**
    $B = B \bigcup b$;
end
$remove\_duplicates(B)$

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 $\sqrt[n_a]{6400000} = 50$. 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

$n = 0$;
$S = \varnothing$;
1 **while** $n < n_{bc}$ **do**
2    $s = combine\_bids\,(A, n_a, B, U, Q)$;
    if $s \neq \varnothing$ **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

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

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

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

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:

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

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

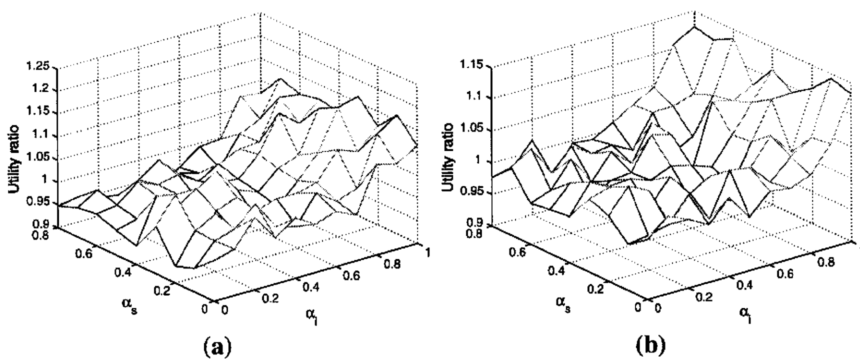

Fig. 6 Individual optimal strategy analysis against symmetric strategy combinations. **a** SA-Q, **b** MWIS-Q

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

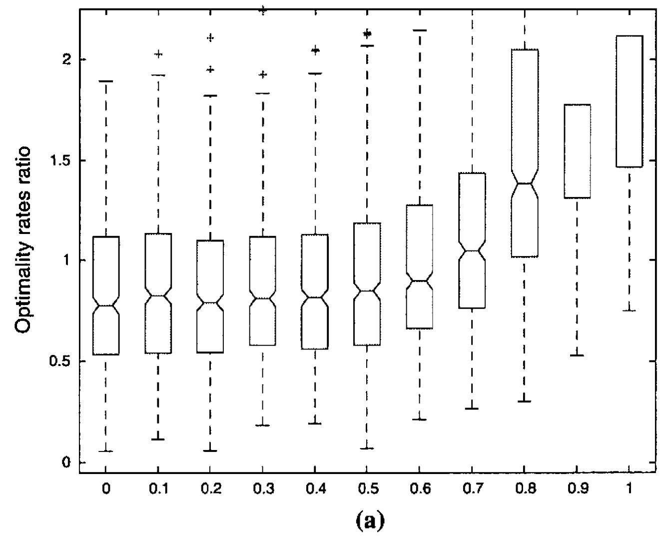

Fig. 7 Individual optimal strategy analysis against random strategy combinations. **a**. SA-Q. **b**. MWIS-Q

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 acount, 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.

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

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 $E\left[sw(s)\right]$ 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:

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

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 al 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 significative 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

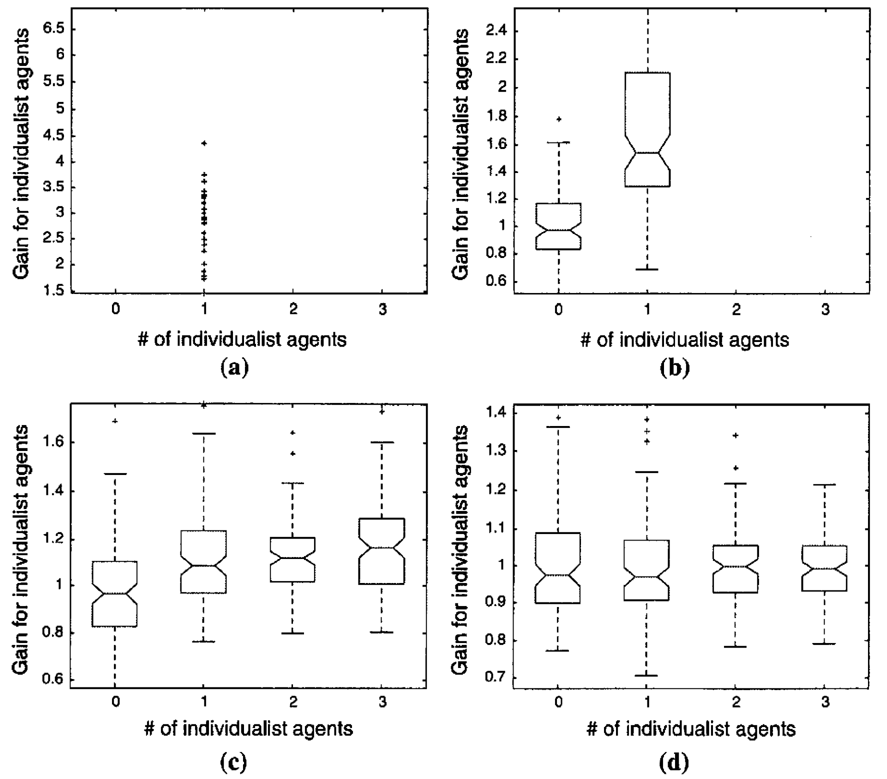

Fig. 8 Stability analysis of the protocol using SA-Q for scenarios with different correlation lengths. **a** $\psi = 2.8$, **b** $\psi = 4$, **c** $\psi = 4.3$, **d** $\psi = 5.9$

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_Q_0*: 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

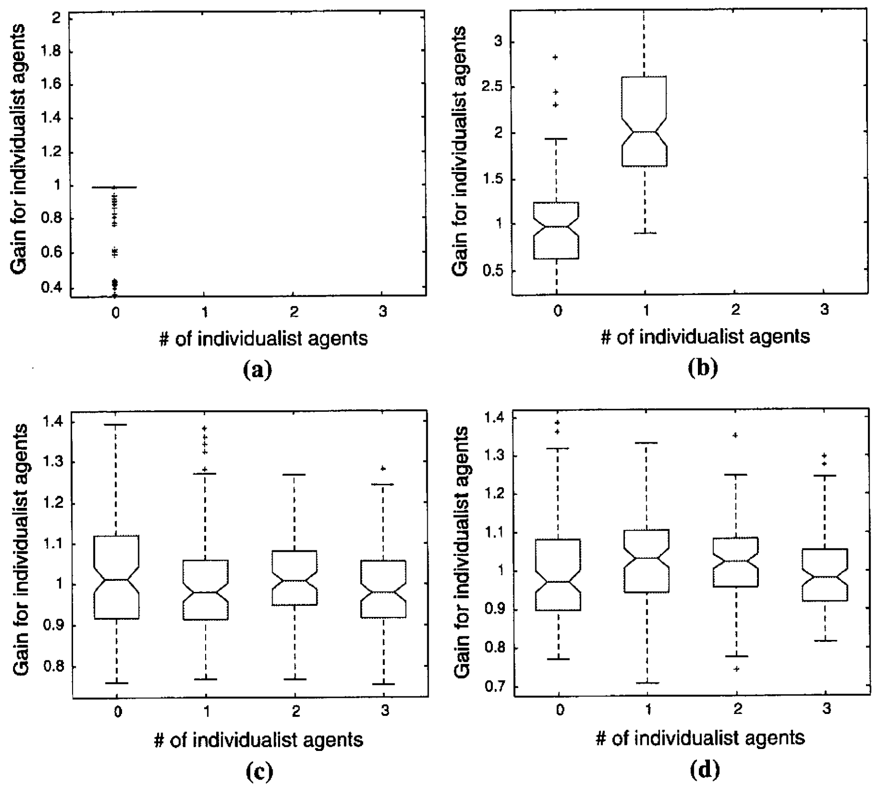

Fig. 9 Stability analysis of the protocol using MWIS-Q for scenarios with different correlation lengths. **a** $\psi = 2.8$, **b** $\psi = 4$, **c** $\psi = 4.3$, **d** $\psi = 5.9$

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]

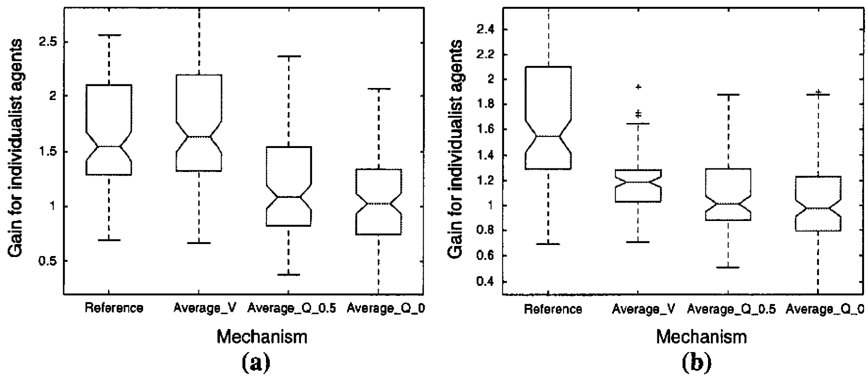

Fig. 10 Effect of the different mechanisms on the stability of the protocol for SA-Q in the most critical scenarios. **a** $\psi = 4.0$, **b** $\psi = 4.3$

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,

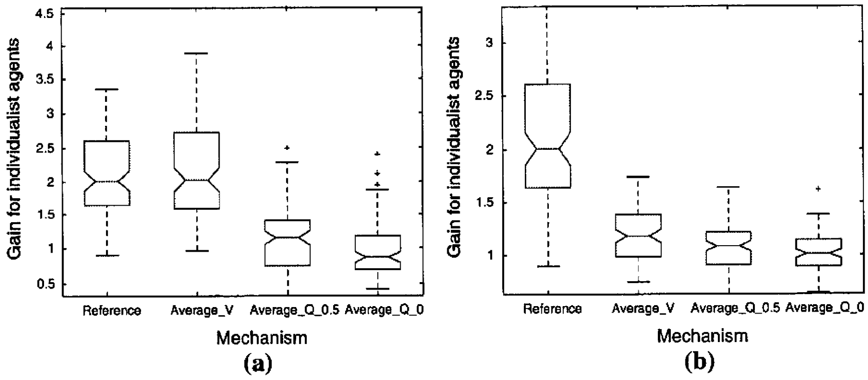

Fig. 11 Effect of the different mechanisms on the stability of the protocol for MWIS-Q in the most critical scenarios. **a** $\psi = 4.0$, **b** $\psi = 4.3$

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.

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

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 $E\left[sw(s)\right]$ 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 eigth 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