

An empirical study of interest-based negotiation

Philippe Pasquier · Ramon Hollands · Iyad Rahwan ·
Frank Dignum · Liz Sonenberg

The Author(s) 2010

Abstract While argumentation-based negotiation has been accepted as a promising alternative to game-theoretic or heuristic-based negotiation, no evidence has been provided to confirm this theoretical advantage. We propose a model of bilateral negotiation extending a simple monotonic concession protocol by allowing the agents to exchange information about their underlying interests and possible alternatives to achieve them during the negotiation. We present an empirical study that demonstrates (through simulation) the advantages of this interest-based negotiation approach over the more classic monotonic concession approach to negotiation.

Keywords Multi-agents systems · Automated negotiation · Interest-based negotiation · Empirical study

This is a substantially revised and expanded version of a paper that appeared in the Proceedings of the Ninth International Conference on Electronic Commerce (ICEC) [39].

P. Pasquier (✉)
School of Interactive Arts and Technology, Simon Fraser University, Surrey, Canada
e-mail: pasquier@sfu.ca

R. Hollands · F. Dignum
Department of Information and Computing Science, Utrecht University, Utrecht, The Netherlands

I. Rahwan
Masdar Institute of Science & Technology, Abu Dhabi, UAE

I. Rahwan
Massachusetts Institute of Technology, Cambridge, USA

I. Rahwan
University of Edinburgh, Edinburgh, UK

L. Sonenberg
Department of Information Systems, University of Melbourne, Melbourne, Australia

1 Introduction

Negotiation is the search for agreement about the exchange (or allocation) of scarce resources among (self-)interested parties. Approaches to one-to-one¹ automated negotiation have been classified in three categories [20]: (1) game theoretic (2) heuristic and (3) argumentation-based.

The first two families focus on the traditional form of automated bilateral negotiation characterized by the exchange of offers between parties with conflicting positions and interests, a style commonly referred to as *position-based negotiation*. These approaches tend to view the object of the negotiation as fixed and reduce the negotiation process to a search problem in the space of possible deals. *Bargaining* consists in an exchange of offers by the agents, who try to accommodate each other's preferences until a deal is acceptable to both parties or the negotiation terminates unsuccessfully.

Argumentation-based negotiation (ABN) has been introduced to enhance automated negotiation with the exchange of richer information between negotiators, supporting or attacking their positions and potentially modifying these positions. Interest-based negotiation (IBN) is a particular type of ABN where the agents exchange information about the goals that motivate their negotiation. The intuition behind ABN (emanating from the realm of human negotiation) is that those exchanges can change the agents' positions in a way that can increase the likelihood or the quality of potential agreements.

While in the last decade, ABN and IBN have been the focus of many publications, as yet very few (if any) empirical evaluations have been provided [44].² IBN advocates the idea that parties can increase the likelihood and quality of an agreement by exchanging information about their underlying goals and about alternative ways to achieve them, and thus influence the agents' preferences over the object of the negotiation [38]. However, no evidence has been provided to support this intuition. The present work advances the state of the art in automated negotiation by testing this hypothesis empirically.

To this end, we define a negotiation model suited to agents with hierarchical goals (Sect. 3). This model includes a bargaining protocol and monotonic concession strategy (Sect. 4.1), a recursive reframing protocol and strategy (Sect. 4.2) and a meta-strategy to articulate those two strategies in the agents' behavior (Sect. 4.3). We exemplify this model and present the simulation tool implemented (Sect. 5). Finally, we compare and discuss the results of (1) negotiations between agents using only the bargaining strategy and (2) negotiations between agents that use both bargaining and reframing, in three different types of encounters (Sect. 6). The next section provides further rationale for the proposed approach by relating it to previous work on automated negotiation.

¹ Many-to-many and many-to-one automated negotiations are usually handled using a growing variety of auction-based models [54,52] and these negotiation types will not be considered in this work.

² A notable exception is the work of Karunatillake et al. [23,22], but it aims at evaluating *social* rather than *interest-based* arguments.

2 Motivations, background and related work

2.1 Position-based negotiation: bilateral bargaining protocols

2.1.1 Game theoretic approaches

Bargaining has received considerable attention in economics, particularly using the analytical methods of game theory [36]. Initial studies dealt with settings where agents have complete information³ about each other's parameters [35], possibly taking into account the effect of additional parameters, such as time preferences and deadlines, on the negotiation outcome [48]. In such cases, one can analytically calculate the (possibly optimal) equilibrium outcome before the bargaining game is played.

Recognising that the assumption of *complete information* is unrealistic in many situations, economists explored bargaining models with *incomplete information*. Complete information, in game theory, means: (1) *common knowledge*, which states that agents have common knowledge about aspects of the game (e.g., each other's private information) and (2) *certain knowledge*, that is no uncertainty nor imprecision. These assumptions are often not satisfied in computational settings that relate to realistic situations, which makes optimal game-theoretic strategies difficult or impossible to prescribe and/or compute. There are at least two ways of dealing with incomplete information in game theoretic approaches.

One approach taken by game theorists is to identify incomplete information with *stochastic information*.⁴ It is then assumed that the world is stochastic (and stationary) and that common information is available, at least in the form of *common prior* probability distributions [17, 49]. Complete probability distributions over the possible values for the information needed are known to the agents. Players could be uncertain about various aspects of other players, such as their discounting factors [49], reservation prices [12], or deadlines [53].

The second approach to deal with incomplete information is *revelation*. When information (about the negotiation object or the opponent's type, goal, reservation price or preferences, ...) is unknown to some agents, it is assumed that a revelation phase will precede the negotiation phase. In this phase, the unknown information will be revealed. The negotiation will then proceed as in the complete information case, under the assumption that the agents are sincere. Sometimes, strategic manipulations can occur at this level [47].⁵ While it allows keeping things analytically tractable in game theory through the use of the revelation principle—revealing all the preferences information—is associated with a number of drawbacks.

Firstly, the revelation of all the information can be computationally very expensive [6]. Indeed, the revelation phase consists of all agents revealing all the unshared information, which may entail too many communications. This revelation phase clearly dissolves the inherent benefit and realism of distributed systems (such as multi-agent systems) in which the information is (and has to stay) distributed [55].

³ In game theory, complete information implies that every player knows the payoffs and strategies available to other players. The notion of complete information should not be confused with notion of perfect information (state of complete knowledge of the actions of the other players associated with most sequential games).

⁴ The underlying idea is that, when the possible values of the information are known, ignorance can often be reduced to equiprobability. However, it is important to notice that this procedure does not really acknowledge the cases where some of the information is not available (e.g., the possible values are not all known).

⁵ Through the concepts of mechanism design, in some special (restricted) cases, game theory allows to ensure sincerity. This happens when (1) the agents are perfectly rational and (2) an incentive compatible mechanism is provided.

Secondly, one of the basics of the principled negotiation approaches [19,5] is that it is important that the agents minimise the amount of information they reveal about their preferences concerning the negotiation object, since any such revelation can weaken their positions [40,46]. For example, negotiators are reticent to reveal their reservation price.

Furthermore, it has been shown that humans minimise the amount of private information they reveal during a negotiation [18]. Any model of automated negotiation should follow the same broad tenet if one wants the agents to be more believable, or simply if one wants to model human negotiation. In particular, human-agent negotiation is a growing field where such considerations are important [16,14,29]. This minimality also serves the purpose of reducing the complexity of communications induced by the revelation of preferences as used in game theoretic approaches.

The idea underlying these two treatments of incompleteness in game-theoretical approaches is to reduce incompleteness and return to a case of complete (possibly stochastic) information.

Underlying most game-theoretic frameworks are a number of other strong assumptions which are often not satisfiable in computational models of bargaining. For example, it is often assumed that each player has unlimited computational resources and time. This implies that each player is capable of computing optimal decisions (e.g., choose the best among alternative offers, compute an optimal strategy,...) and that such computations are performed without cost. This is an unrealistic assumption both in human and software agent systems. Indeed, most game-theoretic models provide no algorithms for implementing such players.⁶

2.1.2 Heuristic approaches

For the above reasons, various *heuristic*-based frameworks for bilateral negotiation have been developed and studied [25]. In these studies, authors typically devise heuristic strategies, or rules-of-thumb, that may be used to produce good (but not necessarily optimal) decisions during negotiation. These strategies are then tested experimentally in order to assess their applicability and performance. A number of models of bilateral bargaining using heuristics have been developed to date, such as:

- *Approaches based on machine learning*: For example, Matos and Sierra [33] proposed a case-based and fuzzy logic based strategy. Zeng and Sycara [57] study bilateral bargaining over a single issue (price), where an agent forms a hypothesis about its opponent's reservation prices (i.e., the highest price he is willing to pay if he is a buyer, or the lowest price he is willing to receive if he is a seller) represented in terms of conditional probability statements (e.g., the probability of an agent's reservation price being 130 is 17%). An agent observes its opponent's offers and uses Bayes' rule to update the estimated opponent reservation price in light of such observations.
- *Constraint-based approaches*: a number of heuristics based on constraint reasoning for offer evaluation and generation have been proposed. A variety of techniques—constraint satisfaction problem (CSP) [1,24], fuzzy constraint satisfaction problem (FCSP) [10], prioritised constraint satisfaction problem (PFCSP) [32]—have been exploited for that

⁶ It is worth mentioning that game theorists have recently realised the significance of computational limitations, resulting in the growth of so-called computational economics. This new field is characterized by its consideration of the work on *bounded rationality* [50]. For example, such models have investigated the impact on the negotiation outcome of limited memory or lookahead in repeated games. In multiagent systems, some work [27] did investigate a single-offer bargaining protocol where computational actions are treated as part of the agent strategy.

purpose. For example, in Faratin et al.'s model [10], a heuristic based on fuzzy similarities selects, among the candidate offers, the one that is most similar to the last offer made by the opponent. The main advantage of constraint-based approaches rests on the expressivity of constraint. Indeed, while regular offers correspond to single points of the search space, constraints usually describe and denote areas of the search space.

- *Approaches based on qualitative decision making*: a number of approaches have exploited various qualitative decision making theories in order to develop models of negotiation addressing cases that are not covered by game-theoretic approaches. These approaches ranges from purely qualitative, like the work of Governatori et al. [15] that rests on defeasible logic, to hybrid approaches that mixes qualitative methods with other heuristics. For example, Jonker et al. [21] developed agents capable of multi-attribute negotiation in the context of incomplete preference information and incorporating “guessing” heuristic, by which an agent uses the history of the opponents bids to predict his preferences. Similarly, the work of Lin et al. [30] tackles the problem of multi-attribute negotiation with bounded rational agents in cases of incomplete information. Validated against human subjects, experiments showed that the resulting agent is performing better than human subjects when playing their role.

All these heuristic strategies make use of the *offer-based information* which is exchanged between agents (e.g., the opponent's rate of concession) to infer useful information (e.g., the opponent's reservation price) which may lead to better agreements.

In summary, heuristic approaches to bilateral bargaining make use of computationally tractable heuristics in order to circumvent the difficulties posed by the high computational demands and the unlimited availability and quality of information required for optimal negotiation behavior.

2.1.3 Summary of the criticisms of position-based approaches

While interesting conceptually, game theoretic approaches are restricted by strong assumptions which are the price paid for enabling formal results and guarantees to be established analytically. More specifically, and putting aside issues of bounded rationality, game theoretic approaches to automated negotiation have been criticized for assuming:

- 1 *Complete and common information*: All the information about the negotiation objects and domain as well as the other negotiators is assumed to be available (at least in the form of probability distributions) and common to all agents. However, “ordinary experience seems to indicate that what makes horse races is variation among prior” [31]. In particular, knowledge of the other agent's utility function is assumed, while “typically, the parties do not know each other's utility functions with any degree of accuracy... Usually, they don't know each other's BATNA [Best Alternative to a Negotiated Agreement]” [56, p. 5].
- 2 *Perfect and correct information*: All the information available (whether it is stochastic or not) is assumed to be accurate. So far, no approach has considered the case in which agents have incorrect⁷ information about the other agents or the negotiation object.

Under those assumptions, game theory can be used either to (1) compute optimal or equilibrium strategies that the agents can use (under the assumption of perfect rationality) or (2) design mechanisms that ensure good properties of the outcome. However, most real world problems are cases of imperfect, erroneous and incomplete (i.e., non-stochastic, not common)

⁷ Not to be confused with uncertain or imprecise.

information where revelation is not realistic (either for computational or strategic reasons, as discussed above).

While allowing progress on those restrictions of game theoretical approaches, heuristic approaches suffer from two other main limitations that they share with game theoretic ones [20]:

- 3 *Agent communication and cognitive capabilities are underused*: The only feedback given to unacceptable proposals is either a counter-proposal or a rejection. The last two decades of work on cognitive agent modelling and multi-agent systems have focused on developing techniques to model goal oriented cognitive agents and the way they influence each other's mental states through communication [55]. It therefore makes sense to exploit these representations when attempting to reach agreement.
- 4 *The positions of the agents are statically defined*: This is the bottom line of position-based approaches to automated negotiation. Each agent has a clearly defined position that is static. The overlap between the agents' positions characterises a fixed negotiation set (i.e., set of possible deals, possibly empty). As noted by Jennings et al. [20], the negotiation space does not change dynamically.

2.2 From position-based negotiation to interested-based negotiation

In this paper, we turn to the idea of *argument-based negotiation* (ABN) [25,26,37], which can also be seen as a special case of heuristic approaches. However, instead of devising new heuristics for making use of the offer-based information exchanged between agents, we follow the intuition that agents may exchange explicit meta-information to improve the way they negotiate. Argumentative messages relevant to the agents' positions are exchanged that can reveal any unknown, non-shared, incomplete, imprecise information about their underlying mental attitudes [44]. *Interest-based negotiation* (IBN) is a particular type of ABN in which the agents can exchange information about their underlying goals to guide the negotiation process.

Originally developed for human negotiation and mediation practices and first introduced by Fisher and Ury [11], the theory of *interest-based negotiation* (IBN) assumes that parties are much more likely to come to a mutually satisfactory outcome when the object of the negotiation is not considered as central as the agents' underlying interests. By focusing on interests to be satisfied rather than positions to be won, IBN allows the agents to search the space of negotiation objects (rather than the space of deals for a particular exchange). When successful, this strategy gives each side more, thereby producing a "win-win" outcome.

While it is still under investigation in the literature on human negotiation [34], IBN has been adapted to negotiation dialogues between artificial agents [41]. In that context, it is a subclass of argumentation-based negotiation (where the agents may argue about other negotiation related issues, i.e., beliefs, goals or social aspects). Interest-based negotiation is based on the idea that the agents can make the goals underlying the negotiation explicit and disclose information about alternative ways to achieve these.

While classical position-based approaches (heuristic approaches or game theoretic approaches) focus on processes to accommodate the agents' preferences with respect to a particular static position, in IBN, the agents' positions may change. We now consider some informal examples to illustrate this idea.

First, consider the following bargaining dialogue where the agents' preferences are summarized by Fig. 1, part (i). In this example, the agents fail to reach a deal because it is not possible to accommodate their respective preferences.




Fig. 1 Upper and lower bounds on agents' preferences

B₁: I would like to rent a car for 4 days please.
S₁: I offer you one for \$400.
B₂: No! How about \$150?
S₂: I can let it go for \$300, but that is my last offer.
B₃: Sorry, I can't afford that amount!

Of course, sometimes, agents' preferences overlap (as in Fig. 1, part (ii)) and a deal can be reached using a monotonic-concession bargaining strategy as in the following dialogue:

B₁: I would like to rent a car for 4 days please.
S₁: I offer you one for \$350.
B₂: No! How about \$200?
S₃: Impossible! How about \$275 then?
B₃: That is OK, I accept!

However, it is worth noticing that agents had to make concessions and move away from their preferred outcome in order to reach a deal. The following dialogue gives an example of IBN dialogue where the seller agent (S) asks the buyer to disclose his underlying interest (goal) before making a concession. In this example, the seller knows an alternative way to fulfill the buyer's underlying goal that he considers to be cheaper.

B₁: I would like to rent a car for 4 days please.
S₁: I offer you one for \$400.
B₂: No! How about \$200?
S₂: I can't do that! What do you need a car for?
B₃: I want to drive to Sydney to attend a conference.

- S₃: You can also fly to Sydney! Booking a plane ticket with Quantum airlines will actually be cheaper!*
- B₄: I didn't know flights were so cheap! Can I get one for \$250?*
- S₄: Yes, that is possible.*

This particular negotiation strategy is known in the human negotiation literature as *reframing*. This reframing allowed the buyer to update the way he was (incorrectly) evaluating the cost of some missing resources. This reframing process allows the agents to move from the situation depicted in Fig. 1, part (i) to the more favorable one described by part (iii) of Fig. 1. Not only was a deal reached (qualitative advantage), but assuming that the seller earns a \$100 profit on both the rental of a car for \$400 and a \$250 Quantum ticket, an agreement is reached without any concession and a benefit is made (quantitative advantage).

The model developed in this paper covers this very common buyer-seller example as well as more complex ones in which agents exchange sets of resources (including compensatory payments) in a context where they both have hierarchical goal structures (with several levels of increasingly abstract underlying goals). In other words, we aim to generalize and characterize under which conditions the phenomenon observed in this realistic scenario occurs.

2.3 Summary of contributions

On the one hand, our work distinguishes itself from game-theoretic approaches in the same way as heuristic approaches do—that is, by providing a more realistic computational model of negotiating agents. In our case, agents do not have any knowledge about the partner's utility function (not even a probability distribution), do not know each other's goals and have erroneous estimations of the value of the resources not owned.

On the other hand, our approach distinguishes itself from heuristic approaches by proposing a model of recursive *reframing*. More precisely, reframing is a particular type of interest-based negotiation (IBN) strategy that allows agents to ask about the goal(s) underlying the requested item(s) and allows them to give constructive feedback. In our case, agents will propose alternative plan(s) that are thought to be cheaper which will allow the correction of misestimated costs for some of the resources being negotiated. This elaborated feedback can change the agents' preferences and thus the issues or objects under negotiation. This strategy may thus allow the agents to reach a deal in this new negotiation space. This IBN strategy can be considered as a type of argumentation-based (ABN) strategy resting on deliberative arguments, where the support consists of goals and beliefs and the conclusion is the goal/intention to get the negotiated item(s).

This paper contributes to the state of the art in automated negotiation in two important ways. Firstly, it provides a formal, computational and reproducible model of IBN agents that partially overcomes the four limitations of position-based negotiation indicated in Sect. 2.1.3. In particular, this model considers the case in which the agents may have erroneous estimations of the value of the resources they need. In the previous example, the buyer agent had an erroneous evaluation of how much a plane ticket costs. As these errors are very common in a variety of realistic situations, our aim is to move towards models that are robust to these types of erroneous evaluations by the agents. The error made by the agent can be an overestimation or an underestimation of any of the resources pursued. Our model defines both a bargaining and a reframing protocol.

Secondly, this paper presents the first extensive empirical analysis of interest-based negotiation. The implementation of our model and a simulator will be presented. We present a number of simulation results providing evidence of the advantage of IBN over strictly

position-based negotiation. These results are presented along two dimensions: qualitative (i.e., whether agents reach some deals when bargaining alone is failing) and quantitative (i.e., the gains made in terms of cost of the plan and benefit made by the agents). Furthermore, we show (experimentally) that in buyer-seller scenarios, using bargaining with reframing Pareto-dominates using bargaining alone, and we discuss reasons why this result does not hold in more complex settings.

Note that the use of simulation was an important methodological choice which resulted from the fact that under the precise set of assumptions made in our model, we did not succeed in developing an analytical proof (for example a game-theoretic analysis) that would give us insights into the consequences or properties of the model.

In summary, this paper presents the first extensive empirical analysis of an interest-based negotiation framework. In our model, agents communicate information about each other's goals and alternative ways to achieve them (estimated to be cheaper by one of the party) and use this information to improve the negotiation process. Such a reframing strategy takes advantage of the communication and cognitive capabilities of goal-driven artificial agents which are mostly ignored by traditional formal and empirical approaches to automated negotiation usually grounded in applied mathematics or micro-economics.

3 Agents with hierarchical goals

This section and the next one present the proposed computational model in a very precise way so as to ensure full reproducibility of the results.

3.1 Preliminary definitions

Our model considers agents that need resources to achieve their hierarchical goal(s). The notion of resource used in the model is kept very general, encompassing physical (e.g., ink to print a page,...) as well as abstract (e.g., an agent i doing an action α at time t) elements of the environment.

Definition 1 (*IBN domain*) An IBN domain consists of the following:

- $\mathcal{A} = \{A_1, \dots, A_n\}$ is the set of agents, limited to two in the context of bilateral negotiation;
- $\mathcal{G} = \{G_1, \dots, G_m\}$ is the set of all possible goals;
- $goal_i \in \mathcal{G}$ is the root goal of agent i ;⁸
- $Res = \{r_1, \dots, r_p\}$ is the set of resources;
- $res_i \subseteq Res$ is the set of indivisible resources owned by agent i ;
- $sub \subseteq \mathcal{G} \times 2^{\mathcal{G} \cup Res}$ is the relationship between goals and their decomposition in sub-goals and resources needed to achieve them;⁹
- $val_i : Res \rightarrow \mathbb{R}$ is a function that returns, for each agent, its valuation of its own resources, as well as the *estimated value* of resources not owned by itself.
- $Budget_i \in \mathbb{R}^+$ is a positive number which stands for the amount of money owned by agent i .

⁸ Throughout the paper, we assume that each agent has a single goal. Multiple goals can be expressed by a single root-level goal that has a single possible decomposition.

⁹ Note that this is a relation (not a function), which allows expressing that a goal may be fulfilled by multiple sets of *alternative* sub-goals or resources.

In that model, money is treated as a finite quantity and will be expressed using positive real numbers. The valuation of resources is expressed using money as a unit. In this context, agents maintain preference intervals over the selling values of the resources they own and over the acquisition costs of the resources they do not possess.

In our model, agents preferences for a given resource are represented using a preference interval as exemplified in the example and real world scenario presented in Sect. 2.2. For each resource x in the domain, each agent maintains a least preferred value $prefmin(x)$ and a most preferred value $prefmax(x)$ for which to exchange the resource. For a resource owned, $prefmin(x)$ denotes the least amount the agent is willing to accept in terms of money, and $prefmax$ the maximum it expects to sell the resource for. For a resource not owned, $prefmin$ corresponds to the maximum amount of money an agent is willing to give to acquire the resource while $prefmax$ is then the least amount of money he expects to have to pay for the acquisition. Note that $prefmin$ returns what is called in economics [2] the *seller's reservation price* for the resource owned (that is the minimum price the agent is willing to accept as a seller of this resource) and the *buyer's reservation price* for the resources not owned (that is the maximum price the buyer is willing to pay for the resource). While in theory, $prefmax$ would be unbounded, in a real world scenario it is not, as no agent would propose to sell a resource owned for an infinite amount of money. Conversely, agents will not propose to acquire a resource for free. In the case of both the buyer and seller, $prefmax$ represents what is sometimes called, in economics and finance, the *first offer price*.

Definition 2 (*Preferences*)

- $prefmin_i : Res \rightarrow \mathbb{R}$ is a function that returns the least preferred value for selling or acquiring a particular resource for agent i .
- $prefmax_i : Res \rightarrow \mathbb{R}$ is a function that returns the most preferred value for selling or acquiring a particular resource for agent i .
- Agents being rational, it follows that:
 - $\forall i \in \mathcal{A}, \forall x \in res_i, prefmax_i(x) \geq prefmin_i(x)$
 - $\forall i \in \mathcal{A}, \forall x \notin res_i, prefmax_i(x) \leq prefmin_i(x)$

Part (i) of Fig. 1 exemplifies $prefmin$ and $prefmax$ and preference intervals in the case of the travel agent scenario presented in Sect. 2.2.

Note that the evaluation of the value of the resources not owned can be (and usually is) different from the reservation price of the owner. In the example of Fig. 1, the rental of the car, a resource noted r , can have a value 200 for the seller S ($val_S(r) = 200$), and can be estimated to 200 by B ($val_B(r) = 200$). Still it is possible that S does not want to sell it for less than 300 ($prefmin_S(r) = 300$, the least amount of money that S is willing to accept for the resource: that is S 's reservation price). It is also possible that B is willing to give up to 250 for acquiring r ($prefmin_B(r) = 250$, the least preferred amount of money B is willing to give to S to acquire r). Note that in that case $prefmin_S(r) - prefmin_B(r) < 0$, indicating that the bargaining interval is empty.

While our model is general and can be used to represent all sorts of scenarios, we further constrain it, for the sake of the empirical studies presented in this paper. We refine the model with the following assumptions that allow us to calculate the preference interval boundaries $prefmin(x)$ and $prefmax(x)$ in a systematic way. First, we make the assumption that the reservation price for a resource ($prefmin$) corresponds to its subjective valuation.

Assumption 1 (*Exchange preferences*) We assume that the least preferred value for selling a resource owned or for the acquisition of a resource not owned (i.e., the reservation price) is its estimated value. Formally:

$$\forall i \in \mathcal{A}, \forall x \in Res, \text{prefmin}_i(x) = \text{val}_i(x)$$

Other choices are possible and are observed in real life scenarios: for example, when a fixed margin is enforced by the seller agent, $\text{prefmin}_i(x) > \text{val}_i(x)$, as is the case in the car rental example of Sect. 2.2 or most products on sales in stores. In other cases, an agent is ready to give up a resource for less than its value, in which case $\text{prefmin}_i(x) < \text{val}_i(x)$. The simplification introduced by Assumption 1 means that: (1) agents may be willing to give up profit¹⁰ on the resources they own but are not ready to do or give anything for less than its subjective value; and symmetrically, (2) agents will not give more than their current valuation for the resources they do not own. In other words, agents are ready to concede their gain up to their reservation price and this equals their subjective evaluation. Note that this assumption does not have any effect on the main results of this paper.

For the first offer price (prefmax), we assume a fixed margin of potential benefit for each agent expressed as a percentage of its reservation price (prefmin). This assumption is introduced to facilitate the presentation of the results and means that for a given agent the absolute ratio between $\text{prefmin}_i(x)$ and $\text{prefmax}_i(x)$ will be identical for all resources.

Assumption 2 (*Potential benefit*) The so-called *potential benefit* of agent i is specified as a percentage b such that:

- $\forall x \in \text{res}_i, \text{prefmax}_i(x) = \text{prefmin}_i(x) + b * \text{prefmin}_i(x)$
- $\forall x \notin \text{res}_i, \text{prefmax}_i(x) = \text{prefmin}_i(x) - b * \text{prefmin}_i(x)$

For example, if the valuation of a particular resource is 100 and the percentage is 10% then the agent will try to sell the resource for 110 (or try to buy it for 90 if she is a buyer). Note that we model b as a percentage mainly for the ease of computing the preference intervals. Other choices could be made without consequences for the main results of the paper.

Another set of assumptions constrains the structure of the previously described IBN domains.

Assumption 3 (*Distribution of the resources*) For distribution of the resources, we assume that:

- $\bigcap_{i \in \mathcal{A}} \text{res}_i = \emptyset$ (the resources are not shared);
- $\bigcup_{i \in \mathcal{A}} \text{res}_i = Res$ (all the resources are owned);

Assumption 4 (*Type of resources*) We assume that resources are consumable.

This will influence the calculations in the sense that—a priori—the estimated cost of using a resource is represented by its subjective value (as given by the function val). Note that—a posteriori—the cost of using a resource is given by its value if it is owned (which can derive from its acquisition cost, but not necessarily) and by its acquisition cost if it has been acquired. For example, in the car rental example of Sect. 2.2, Fig. 1 ii, agent B evaluates the use of a rented car for four days to \$300 ($\text{val}_B(4\text{daysCarRental}) = 300$). Using that resource as part of a plan will thus cost B \$300 (i.e., B 's assets will diminish by \$300). This is if B already paid \$300 for the car rental. If B does not own that resource yet and merely plans to acquire it, this is just an estimated value and she may end up getting a deal for \$275 (as in the example of Figure 1 ii) and will then have to update her valuation consequently. In any case, considering reusable resources would complexify that calculus and is left as future work.

¹⁰ Profit is the difference between the actual value for which a resource is exchanged and the subjective evaluation of that resource prior to the negotiation.

Assumption 5 (*Shared vs. private knowledge*) We assume that the agents have a shared, common, and accurate knowledge of the set of all possible goals, the set of all possible resources, and all possible decompositions of goals, i.e., they know \mathcal{G} , Res , res_i and sub . In other words, they all have the “know-how” of the domain and have a common ontology thereof. However, they do not know each other’s goals, valuation function or preferences. I.e, $goal_i$, val_i , $prefmin_i$ and $prefmax_i$ are strictly private information.

While the presented model is more general, in the remainder of this paper, we restrict the negotiation to two agents. Exploring the model beyond bilateral negotiation settings is left to future work.

Algorithm 1 Generate plans for goal G

Global variables:

$plan[]$: Vector of plans, i.e. trees
 $z := 0$: Number of plans in $plan[]$
 $cost[]$: Vector of costs for the corresponding plans
 $generatelist[]$: Vector of lists of nodes to generate

Global code:

$z := 1$; $plan[1] := G$: Initialize $plan[1]$ with G
 $generate(G, plan[1])$: Generate all the plans for goal G
 $i := \min(cost[])$: Index of the cheapest plan
 $result := plan[i]$: Return the cheapest plan

Procedure $generate(node, plan[x])$

Local variables:

$i := 1$: Number of the current decomposition for the current node
 $next, child$: Variables of type node, i.e. a goal or a resource
 $children$: Set of nodes
 $curplan := plan[x]$: Current plan
 $curind := x$: Index of the current plan
 $curgeneratelist := generatelist[x]$

Body:

```

for all  $(node, children) \in sub$  do
  if  $i > 1$  then
     $plan[z + 1] := curplan$ 
     $z := z + 1$ 
     $curind := z$ 
  end if
   $i := i + 1$ 
  for all  $child \in children$  do
     $addchild(child, node, plan[curind])$ 
    if  $child \in Res$  then
       $cost[curind] := cost[curind] + val(child)$ 
    else
       $add(child, generatelist[curind])$ 
    end if
  end for
  while  $generatelist[curind] \neq \emptyset$  do
     $next := pop(generatelist[curind])$ 
     $generate(next, plan[curind])$ 
  end while
end for

```

3.2 Generating hierarchical plans

Under the assumptions sketched above, agents can use Algorithm 1 given below to generate all the possible plans (along with their costs) to achieve a particular goal.

Definition 3 (*Plan, cost and benefit*) The n th plan P_i^n generated by agent i for achieving a goal G_0 is a tree such that:

- G_0 is the root;
- Each non-leaf node is a goal $G \in \mathcal{G}$ with children x_1, \dots, x_m such that:
 $sub(G, \{x_1, \dots, x_m\})$;
- Each leaf node x is a resource: $x \in Res$

We denote $needed(P_i^n)$ the set of leaf nodes of the plan P_i^n and $missing_i(P_j^n)$ the subset of $needed(P_j^n)$ not owned by agent i ($missing_i(P_j^n) = needed(P_j^n) \setminus res_i$). More generally, we denote by $missing_i$ the set of resources that an agent i is willing to acquire.

The estimated cost of a plan P_j^n for agent i is:¹¹

$$cost_i(P_j^n) = \sum_{x \in needed(P_j^n)} val_i(x)$$

The potential benefit that an agent i can make on a plan P_j^n —in selling the resources he owns that are needed by agent j for executing the plan—is defined by:

$$benef_i(P_j^n) = \sum_{x \in needed(P_j^n) \cap res_i} prefmax_i(x) - prefmin_i(x)$$

We note $>_i^c$ and $>_i^b$ the preference ordering of the plans according to these estimated costs and benefits.

The use of the hierarchical structure of goals, super-goals and subgoals gives a great expressivity to the model [45]. Since we do not over-specify this choice, related representations using the same type of structures (like tasks, sub-tasks and super-tasks, as in TAEMS [28]) can still be captured.

Assumption 6 (*Independence of goals and plans*) We assume that no overlap exists between agents' needed resources, nor between their plans' root or sub-goals:

- $\forall n, m, needed(P_i^n) \cap needed(P_j^m) = \emptyset$
- $\forall n, m, \forall x \in \mathcal{G} \cup Res : x \in P_i^n \Rightarrow x \notin P_j^m$

This last assumption is realistic in many domains where agents operate in separate but complementary sub-domains. For example, agents that are associated with different roles in an organization often have mutually exclusive sets of goals. Other examples include most of the buyer-seller cases (that constitutes the bulk of economic exchanges). This assumption means that there will not be positive or negative interactions between goals, nor conflicts over resources. These issues are dealt with in related work [47], including ours [43], and we plan to address them further in future work (described in Sect. 8).

Despite these assumptions, it is important to notice that this model is still very general, encompassing multiple instances. It moves beyond usual position-based approaches by

¹¹ Note that we use “cost” for plans and “valuation” for resources. Indeed, we assume that the resources will be consumed by the plan execution as specified by Assumption 4. Note also that i and j can differ as an agent can estimate the cost of another agent’s plan.

addressing the case in which agents build their own (i.e., subjective and possibly erroneous) evaluations of resources according to their own calculus. While the acquisition of these evaluations is not modelled in the present paper, these can be based on the agents' beliefs (possibly uncertain, imprecise or erroneous), biased by their interests as well as their past experience, including information that has been communicated to them by other agents and so on.

It is usually the case that these subjective evaluations are erroneous and agents cannot (easily) assess how far they are from the real value since: if an agent asks the opponent about the true value, the latter will (rationally) give her preferred value. Position-based negotiation (called bargaining throughout this paper) is the process by which the agent will try to move the partner's position from her preferred value toward her reservation value (reservation price in the rest of this paper) which can itself be different from the subjective value given by the owner of the resources.

These assumptions are thus realistic modelling choices, that correspond to real world situations such as price negotiation on a vegetable market. The buyer typically estimates the value of the resources he needs prior to the bargaining. Anyone who has travelled will know how erroneous such estimations can be. Not only are these cases interesting to model, but they are also relevant to automated negotiation for the reasons mentioned above. A recent synthesis on automated negotiation [3] indicates the need for models that deal with uncertainty or imprecision about the negotiation object rather than uncertainty about the partner utility function (or type) as is usually the case.

4 The negotiation framework

In order to enable agents to use both bargaining and reframing, one needs to define appropriate communication protocols and negotiation strategies. Figure 2 presents the UML 2.0 specification of the two sub-protocols: (a) the bargaining protocol and (b) the reframing one. The following subsections describe these protocols as well as the associated strategies. A detailed example of negotiation using this framework is presented in Sect. 5.2.

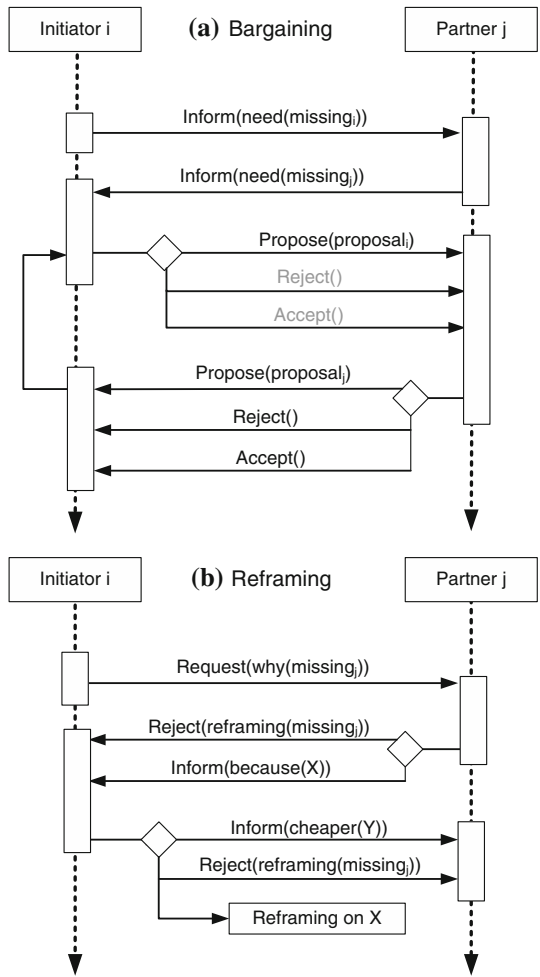
4.1 Bargaining: protocol and strategy

Part (a) of Fig. 2 describes the proposed bargaining protocol. In order to ensure that at least the initiator agent needs a negotiation dialogue to occur, we assume that $missing_i \neq \emptyset$. The bargaining protocol initiated by agent i with agent j is divided into two parts as follows:

- **Part one: negotiating the resources to be exchanged**
Each agent discloses the set of resources that he wants:¹²
 $inform_i(need(missing_i)), inform_j(need(missing_j))$
- **Part two: bargaining over the payment**
 - i makes a first offer (Definition 7)
 - Then j chooses between the three following options:
 - *accept i 's proposal*: this option is chosen by an agent if the condition for the acceptance of a proposal (Definition 6) is met, in which case bargaining terminates with a deal;

¹² Assumption 6 simplifies this part of the protocol. Relaxing it to treat negative and positive interactions between goals and needs is left as future work.

Fig. 2 UML 2.0 specification of the bargaining and reframing protocols



- *reject i's proposal*: this option is chosen if the ending condition (Definition 9) is met, in which case bargaining terminates without a deal;
- *make a counter proposal*: a counter proposal is generated according to the bargaining strategy (Definition 8 or 7 if it is *j*'s first proposal), in which case the negotiation partner has to respond similarly, with acceptance, rejection or counter proposal.

Note that in part two of the protocol, the set of resources negotiated cannot be changed anymore (*missing_i* and *missing_j* are fixed); that is the bargained items are fixed (in conformance with classic definitions of bargaining in economics).

Definition 4 (Proposal) A proposal (or offer) from *i* to *j* is a tuple: $\langle S_{wanted}, S_{given}, Payment \rangle_{i \rightarrow j}$, where *S_{wanted}* is a set of resources wanted by *i* from *j*, *S_{given}* is a set of resources given to *j* and *Payment* is an amount of money offered by *i* to *j* (if positive) or asked (if negative) to compensate any difference in value (which is

subjectively evaluated by i). We note p_i^t , the t -th proposal issued by agent i and abbreviate it to p_i when it does not matter.

Definition 5 (Proposal evaluation) Given a proposal $p_i = \langle S_{wanted}, S_{given}, Payment \rangle_{i \rightarrow j}$, its subjective evaluation by agents i and j respectively is defined as follow:

$$eval_i(p_i) = \sum_{x \in S_{wanted}} val_i(x) - \sum_{x \in S_{given}} val_i(x) - Payment$$

$$eval_j(p_i) = Payment + \sum_{x \in S_{given}} val_j(x) - \sum_{x \in S_{wanted}} val_j(x)$$

Definition 6 (Acceptance of a proposal) An agent j will accept a proposal $p_i = \langle S_{wanted}, S_{given}, Payment \rangle_{i \rightarrow j}$ iff:

- $eval_j(p_j) - eval_j(p_i) \leq \rho$, where p_j is the next proposal to be issued by j and ρ is a strictly positive real number standing for the deviation tolerated by j .¹³
- $Payment \leq Budget_j$

We assume that agreements are enforceable.

Definition 7 (Bargaining initial proposals) For an agent j , the first offer will take the classic form:

$$p_j^1 = \langle missing_j, missing_i, Payment \rangle_{j \rightarrow i}$$

where $Payment$ is defined as:

$$Payment = \min \left(\sum_{r_i \in missing_j} prefmax_j(r_i) - \sum_{r_j \in missing_i} prefmax_j(r_j), Budget_j \right)$$

In this paper, we will use a simple monotonic concession strategy.

Definition 8 (Concession strategy) Given a proposal $p_i = \langle S_{wanted}, S_{given}, Payment \rangle_{i \rightarrow j}$ received by j as a response to his previous proposal p_j^t , the next proposal to be issued by j would take the form:

$$p_j^{t+1} = \langle S'_{wanted}, S'_{given}, Payment' \rangle_{j \rightarrow i}$$

where:

- $S'_{wanted} = S_{given}$;
- $S'_{given} = S_{wanted}$;
- $Payment'$ is such that:

$$eval_j(p_j^{t+1}) = \frac{(eval_j(p_j^t) + eval_j(p_i))}{2}$$

If $eval_j(p_j^{t+1}) \geq 0$ and $Budget_j \geq Payment'$ then:

$$p_j^{t+1} = \langle S'_{wanted}, S'_{given}, Payment' \rangle_{j \rightarrow i} \text{ or else } p_j^{t+1} = p_j^t.$$

¹³ ρ is usually quite small and is here just to avoid infinite length bargaining. It can be expressed as a percentage of $eval_j(p_j)$, typically between 0.1 and 5% depending on the domain.

The last part of this definition implies that when an agent reaches a point where he cannot make any more concessions (whether because he reached his least preferred acceptable proposal or because he does not have enough money), he will repeat his last proposal. When both agents are in that situation, the bargaining ends without a deal as specified by the following ending condition.

Definition 9 (*Bargaining ending condition*) The bargaining ending condition is reached iff i 's two last proposal p_i^t and p_i^{t+1} are such that $p_i^t = p_i^{t+1}$ and j 's last and forthcoming proposal p_j^t and p_j^{t+1} such that $p_j^t = p_j^{t+1}$.

In that case, j will issue a *reject* message rather than the proposal p_j^{t+1} .

4.2 Reframing: protocol and strategy

Part (b) of Fig. 2, describes the (recursive) reframing protocol. Initiated by agent i , this protocol allows agent i to ask agent j what is(are) his underlying goal(s) justifying his need for $missing_j$. Agent j can then (1) disclose his set of underlying goals $g_j = \{x_1, \dots, x_n\}$ motivating the need to acquire $missing_j$ while clarifying the sub-plan(s) selected for achieving it¹⁴ or (2) reject the question, thus ending the reframing protocol (e.g., there is no underlying goal to disclose, the last goal disclosed was the root goal).

Agent i then generates all the possible plans for achieving the goals of g_j and can either:

1. Inform j of **one** alternative plan to achieve one of j 's underlying goals that i believes (according to \succ_i^c built upon i 's private valuation function val_i) to be preferable (that is cheaper) than the one previously selected by j .
2. Reject the reframing (e.g., because he does not have any alternative offer to propose for any of the goals pursued by j , nor for any of their potential super-goals).
3. Start a new reframing protocol to inquire about the super-goals of g_j .

In the first case, when several alternative ways to achieve one of the goals of g_j exist, revelations are made by agent i in a rational way according to \succ_i^b : that is, the one with the highest potential benefit for i is disclosed first.¹⁵

On receiving the information that there is an alternative plan (P_j^2) for achieving a goal from g_j which is evaluated by i to be cheaper than the one selected by j (P_j^1), agent j will update his valuation function over the resources not owned according to this new information. We assume that the valuation of the resources owned is fixed and it will not be updated. In general, there are many ways in which this update can be made. When possible, j can (1) raise P_j^1 's cost, (2) lower P_j^2 's cost or (3) both (and to various degrees). In the current implementation, we use an update strategy of the type (3).

Definition 10 (*Update function*) Let $missing_j(P_j^1)$ and $missing_j(P_j^2)$ be the sets of resources not owned by j involved in P_j^1 (j 's current plan) and P_j^2 (alternative plan proposed by i) respectively:

- The values of the resources that are shared by the two plans (i.e. included in: $missing_j(P_j^1) \cap missing_j(P_j^2)$) are not changed.

¹⁴ Only the goals of one level up are revealed.

¹⁵ In case of equal potential benefit, a random choice is made.

- The values of the resources that are not shared by the two plans, i.e. included in: $missing_j(P_j^1) \setminus (missing_j(P_j^1) \cap missing_j(P_j^2))$ or in $missing_j(P_j^2) \setminus (missing_j(P_j^1) \cap missing_j(P_j^2))$ are equally raised and lowered respectively so that the cost of the sub-plans are such that $cost_j(P_j^2) = cost_j(P_j^1) - \rho$ (i.e. P_j^2 is cheaper than P_j^1).¹⁶

Other choices corresponding to strategies of type (1) or (2) are worth studying but are left as future work. In any case, these update strategies assume that agent j trusts i since the update makes i 's statement true in j 's model.¹⁷

4.3 Agent behavioural model

In this paper, we are mainly interested in comparing the results of negotiations between agents capable of bargaining only (noted *BO*) versus between agents capable of bargaining and reframing (noted *BR*). The following sub-sections describe the execution cycle of these two types of agents.

4.3.1 Bargaining only agents

A BO agent's execution cycle can be summarised as follows:

1. The agent generates all the possible plans to achieve her goal and orders them according to their costs;
2. She selects the cheapest plan to achieve her goal;
3. If the plan involves resources not owned, then she starts a bargaining as described in Sect. 4.1;
4. If the bargaining fails, she withdraws the current plan and proceeds with the next cheapest plan through step 3.

The process terminates when there is no plan left or when a deal is reached.¹⁸

4.3.2 Reframing capable agents

Bargaining and reframing (BR) capable agents' execution cycle extends the one of BO agents with the reframing capabilities described in Sect. 4.2.

Since BR agents have two different negotiation strategies available to them—namely bargaining and reframing—there are a variety of meta-strategies available to compose them. In particular, if we note by B a complete bargaining, B_1 the first part of the bargaining protocol (i.e. only the revelation of needed resources), R_A a reframing initiated by A and U the fact that the agents' valuation functions are updated (or not) according to the last reframing, the following meta-strategies give different outcomes:

¹⁶ In the absence of more information, it is assumed that the alternative plan is just a little bit cheaper (where "a little bit" is represented by ρ).

¹⁷ This assumption is to be relaxed. However, the assumption rests on the intuition that it is to the advantage of both agents to be sincere and trust each other's statements.

¹⁸ Since each agent has a finite number of plans and the monotonic concession protocol is known to terminate [9], negotiations between BO agents always terminate.

1. $B_1 \rightarrow R_B \rightarrow U \rightarrow B \rightarrow R_A \rightarrow B\dots$
2. $B_1 \rightarrow R_A \rightarrow U \rightarrow B \rightarrow R_B \rightarrow B\dots$
3. $B \rightarrow R_A \rightarrow U \rightarrow B \rightarrow R_B \rightarrow U \rightarrow B\dots$
4. $B \rightarrow R_B \rightarrow U \rightarrow B \rightarrow R_A \rightarrow U \rightarrow B\dots$
5. $B \rightarrow R_A \rightarrow U \rightarrow R_B \rightarrow U \rightarrow B\dots$
6. $B \rightarrow R_{A\&B} \rightarrow U \rightarrow B\dots$

The two first strategies start with the reframing protocol (as soon as the resources to be exchanged are known). This can lead to revelations that can modify the agents' preferences before the first bargaining occurs. In addition to all the undesirable properties of using the revelation principle (even partial) described in Sect. 2.1.1, this can have negative side effects on the negotiation outcome. This can be exemplified by looking at a modified version of the informal example given in Sect. 2.2. We assume that: (a) cars are cheaper than plane tickets and (b) for a given price, the seller is making more benefit from a plane ticket than from a car. This meta-strategy could then penalise the buyer. Indeed, an early revelation of the buyer's underlying goal would result in the seller wanting to hide the fact that there are some cars available, or artificially augmenting his reservation price for car rental in the first place to see if he can sell a plane ticket. In general, agents have an incentive not to reveal their underlying goals before they have evidence that their initial preferences cannot be satisfied (i.e., the initial bargaining fails).

In this paper, BR agents will use the meta-strategy number 6. Agents start with a (complete) bargaining (protocol). If the bargaining fails, both agents attempt¹⁹ reframing before initiating a new bargaining. All reframings have to terminate before the agents (possibly) update their valuation functions and a new bargaining is initiated. The rationale for this choice is that such "parallel" reframing does not give any undue advantage to one agent. Note that all other cases create different asymmetries between the agents that can modify substantially the results obtained. Studying these other meta-strategies is left for future work.

To further avoid any asymmetry between the agents (in the BR as in the BO case), the agent that *initiates* the negotiation in our simulations is chosen randomly.

5 Implementation and example

5.1 Implementation and parameters of experimentation

In order to assess the model's behavior, we have developed a negotiation simulator based on the 3APL agent programming language and environment [7,8]. This simulator allows the generation of synthetic IBN domains, as well as BO and BR agents as well as conducting and visualizing various negotiation simulations. A variety of parameters can be manipulated as described below.

5.1.1 Parameters about the domain

It is clear that the structure of the domain can influence the results obtained in the experiments envisioned to compare BO and BR negotiations. In particular, the complexity and richness of the goal decomposition relation (*sub*) will have an impact on the usefulness of using reframing. For example, reframing is completely useless when no alternative decompositions are available whatsoever. In those cases, no alternative ways to achieve the underlying goal(s) can be proposed or discussed.

¹⁹ Following a buyer-seller bargaining only one agent can attempt a reframing.

Our simulation tool includes a synthetic domain generation module which allows defining the depth of the trees (i.e., plans) generated as well as their branching factors. Finally, the number of alternative ways to achieve each goal can be manipulated as well. As an example, with depth 2, branching factor 2 and a number of alternatives of 2, we get 8 possible plans to achieve the agent's main goal.

5.1.2 Parameters related to the agents

For each agent introduced in the system, the resources owned are distinguished from the resources not owned. For each resource, a valuation for the agent who owns it is chosen randomly between 50 and 500. Then, the valuation for the other agent is calculated according to the "error" the agent is making in evaluating the resources he does not own. We use a Gaussian distribution to encode this error, where the mean (noted *error*) of the distribution and standard deviation (noted *stdvar*) are expressed as percentages of the valuation attributed to the resource by the system for its owner. Finally—for both agents—the preference boundaries (*prefmin* and *prefmax*) are calculated for each resource according to the *percentage of potential benefit* over the reservation price that an agent will try to make when buying or selling resources (Assumption 2).

5.2 Detailed example

Let us look at an example to illustrate the model as well as our simulation parameters. In order to clarify the notation, we use $[x, y]_i^j$ as a shorthand for $\text{prefmin}_i(r_j) = x$ and $\text{prefmax}_i(r_j) = y$ when $r_j \in \text{res}_i$ (in that case: $x \leq y$). In the same way, we use $[y, x]_i^j$ as a shorthand for $\text{prefmin}_i(r_j) = x$ and $\text{prefmax}_i(r_j) = y$ when $r_j \notin \text{res}_i$ (in which case: $x \geq y$). Let an IBN domain be such that:

- $\mathcal{A} = \{A, B\}$ is the set of agents;
- $\mathcal{G} = \{G_1, \dots, G_{17}\}$ is the set of all possible goals;
- $\text{goal}_A = G_1$ and $\text{goal}_B = G_8$;
- $\text{Res} = r_1, \dots, r_{20}$;
- $\text{sub} = \{(G_1, \{G_2, G_3\}), (G_1, \{G_4, G_5\}), (G_1, \{G_6, G_7\}), (G_1, \{G_8, G_9\}), (G_2, \{r_1, r_2\}), (G_3, \{r_3, r_4\}), (G_4, \{r_5, r_6\}), (G_5, \{r_7, r_8\}), (G_6, \{r_9, r_{10}\}), (G_7, \{r_{11}, r_{12}\}), (G_8, \{G_9, G_{10}\}), (G_8, \{G_{11}, G_{12}\}), (G_9, \{r_{13}, r_{14}\}), (G_{10}, \{r_{15}, r_{16}\}), (G_{11}, \{r_{17}, r_{18}\}), (G_{12}, \{r_{19}, r_{20}\}), (G_{13}, \{r_1, r_2\}), (G_{13}, \{r_3, r_4\}), (G_{14}, \{r_5, r_6\}), (G_{14}, \{r_7, r_8\}), (G_{15}, \{r_9, r_{10}\}), (G_{15}, \{r_{11}, r_{12}\}), (G_{15}, \{G_2, G_3\}), (G_{16}, \{G_4, G_5\}), (G_{16}, \{G_6, G_7\}), (G_{16}, \{G_8, G_9\}), (G_{17}, \{r_1, r_2\}), (G_{17}, \{r_3, r_4\}), (G_{17}, \{r_5, r_6\}), (G_{17}, \{r_7, r_8\}), (G_{17}, \{r_9, r_{10}\})\}$;
- $\text{res}_A = \{r_1, r_7, r_8, r_9, r_{10}, r_{11}, r_{12}, r_{14}, r_{15}, r_{16}, r_{17}, r_{18}, r_{19}, r_{20}\}$
- $\text{Budget}_A = 2000$
- $\text{res}_B = \{r_2, r_3, r_4, r_5, r_6, r_{13}\}$;
- $\text{Budget}_B = 1500$
- The values of $\text{val}_A, \text{val}_B, \text{prefmin}_A, \text{prefmin}_B, \text{prefmax}_A$ and prefmax_B , have been generated with $\text{error}_A = \text{error}_B = 0\%$ and $\text{stdvar}_A = \text{stdvar}_B = 70\%$ and with potential benefit of 11%:

$$[353, 388]_A^{r_1}, [308, 343]_A^{r_2}, [51, 57]_A^{r_3}, [411, 456]_A^{r_4}, [265, 294]_A^{r_5}, [457, 508]_A^{r_6}, [86, 94]_A^{r_7}, [268, 295]_A^{r_8}, [410, 451]_A^{r_9}, [254, 278]_A^{r_{10}}, [103, 113]_A^{r_{11}}, [433, 476]_A^{r_{12}}, [220, 244]_A^{r_{13}},$$

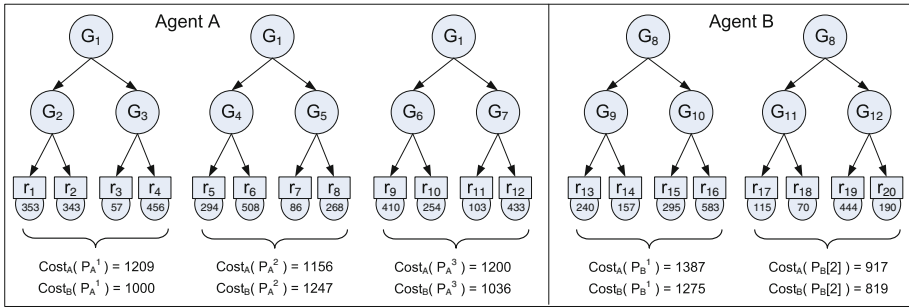


Fig. 3 Plans initially generated by agents A and B

$[371, 408]_A^{r_{14}}, [200, 220]_A^{r_{15}}, [468, 515]_A^{r_{16}}, [114, 126]_A^{r_{17}}, [95, 105]_A^{r_{18}}, [470, 517]_A^{r_{19}}, [154, 169]_A^{r_{20}}$, and
 $[147, 174]_B^{r_1}, [314, 346]_B^{r_2}, [66, 73]_B^{r_3}, [445, 490]_B^{r_4}, [432, 475]_B^{r_5}, [427, 470]_B^{r_6}, [80, 89]_B^{r_7}, [222, 247]_B^{r_8}, [262, 328]_B^{r_9}, [171, 214]_B^{r_{10}}, [72, 90]_B^{r_{11}}, [323, 404]_B^{r_{12}}, [240, 264]_B^{r_{13}}, [126, 157]_B^{r_{14}}, [266, 295]_B^{r_{15}}, [524, 583]_B^{r_{16}}, [103, 115]_B^{r_{17}}, [63, 70]_B^{r_{18}}, [400, 444]_B^{r_{19}}, [171, 190]_B^{r_{20}}$

In this example, we assume that $\rho = 1$ (see Definition 6). The agents generate all the possible plans to achieve their main goals along with their costs as shown in Fig. 3. With those domain values, we will exemplify both the BO and the BR cases.

BO case

According to the proposed bargaining protocol, agents A and B first reveal their needs. The preferred plans according to \succ_A^c and \succ_B^c are P_A^2 and P_B^2 respectively (see Fig. 3). This leads to: $missing_A = \{r_5, r_6\}$ and $missing_B = \{r_{17}, r_{18}, r_{19}, r_{20}\}$. As shown in the following table, the agents start the second part of the bargaining protocol with their preferred offers B^1 and A^1 (Definition 7). Then, following the concession strategy (Definition 8), B repeats his offer and A cannot change her's either, indicating that the ending condition (Definition 9) is reached and the bargaining process fails.

Nb	Message	$eval_A$	$eval_B$
B^1	$\langle \{r_{17}, r_{18}, r_{19}, r_{20}\}, \{r_5, r_6\}, -208 \rangle_{B \rightarrow A}$	-177	+168
A^1	$\langle \{r_5, r_6\}, \{r_{17}, r_{18}, r_{19}, r_{20}\}, -195 \rangle_{A \rightarrow B}$	+164	-235
B^2	$\langle \{r_{17}, r_{18}, r_{19}, r_{20}\}, \{r_5, r_6\}, -208 \rangle_{B \rightarrow A}$	-177	+168
A^2	<i>reject</i>		

After this bargaining, the two agents withdraw their selected plans. For agent A, P_A^3 is the next preferred plan. Because P_A^3 has no missing resource, agent A will achieve her goal on her own. Agent B's next preferred plan is P_B^1 which involves missing resources r_{14}, r_{15} and r_{16} . The following table summarizes the second bargaining that fails:

Nb	Message	$eval_A$	$eval_B$
B^1	$\langle \{r_{14}, r_{15}, r_{16}\}, \{\}, +916 \rangle_{B \rightarrow A}$	-123	+119
A^1	$\langle \{\}, \{r_{14}, r_{15}, r_{16}\}, -1143 \rangle_{A \rightarrow B}$	+104	-108
B^2	$\langle \{r_{14}, r_{15}, r_{16}\}, \{\}, +1029.5 \rangle_{B \rightarrow A}$	-9.5	+5.5
...
A^4	$\langle \{\}, \{r_{14}, r_{15}, r_{16}\}, -1043.7 \rangle_{A \rightarrow B}$	+4.7	-8.7
B^5	$\langle \{r_{14}, r_{15}, r_{16}\}, \{\}, +1029.5 \rangle_{B \rightarrow A}$	-9.5	+5.5
A^5	<i>reject</i>		

In this case, only agent *A* succeeds in achieving her goal with a plan P_A^3 that has a cost of 1200.

BR case

In the BR case, everything is identical to the BO case until the end of the first bargaining, at which point the agents will both try to reframe, leading to the conversation represented in the following table. First, they ask each other for the reasons behind the requested resources. After revealing these, they each compute the possible plans (along with their costs) for the other agent to achieve his/her goal(s) and since they do not have any cheaper alternative to propose, they iterate and start a new reframing. While *A*'s third reframing attempt leads to a rejection (utterances A^5 and B^6), *B* proposes a cheaper plan for *A*'s goal G_1 in his second reframing attempt (utterance B^3 , A^4 and B^5 , see Fig. 3 for the costs). Note that *B* proposes P_A^1 rather than P_A^3 according to $>_B^b$ (P_A^3 does not allow any benefit).

Nb	Perf.	Message
B^1	<i>request</i>	<i>why</i> $\langle\{r_5, r_6\}\rangle_{B \rightarrow A}$
A^1	<i>request</i>	<i>why</i> $\langle\{r_{17}, r_{18}, r_{19}, r_{20}\}\rangle_{A \rightarrow B}$
B^2	<i>inform</i>	<i>because</i> $\langle\{G_{11}, \{r_{17}, r_{18}\}\}, \{G_{12}, \{r_{19}, r_{20}\}\}\rangle_{B \rightarrow A}$
A^2	<i>inform</i>	<i>because</i> $\langle\{G_4, \{r_5, r_6\}\}\rangle_{A \rightarrow B}$
B^3	<i>request</i>	<i>why</i> $\langle\{G_4\}\rangle_{B \rightarrow A}$
A^3	<i>request</i>	<i>why</i> $\langle\{G_{11}, G_{12}\}\rangle_{A \rightarrow B}$
B^4	<i>inform</i>	<i>because</i> $\langle\{G_8, \{r_{17}, r_{18}, r_{19}, r_{20}\}\}\rangle_{B \rightarrow A}$
A^4	<i>inform</i>	<i>because</i> $\langle\{G_1, \{r_5, r_6, r_7, r_8\}\}\rangle_{A \rightarrow B}$
B^5	<i>inform</i>	<i>cheaper</i> $\langle\{G_1, \{r_1, r_2, r_3, r_4\}\}\rangle_{B \rightarrow A}$
A^5	<i>request</i>	<i>why</i> $\langle\{G_8\}\rangle_{A \rightarrow B}$
B^6	<i>reject</i>	<i>why</i> $\langle\{G_8\}\rangle_{B \rightarrow A}$

Agent *A* will update her valuation function with respect to the new information (Definition 10), the values of the resources will be changed, resulting in updated plans' costs. The evaluated cost of P_A^1 for *A* becomes 1191 (with the updated preference intervals $[265, 332]_A^{r_2}$, $[37, 64]_A^{r_3}$, $[356, 445]_A^{r_4}$), and the cost of P_A^2 becomes 1192. Agent *A*'s preferred plan is now P_A^1 . The following table summarizes the next bargaining, in which a deal is reached.

Nb	Message	$eval_A$	$eval_B$
A^1	$\langle\{r_2, r_3, r_4\}, \{r_{14}, r_{15}, r_{16}\}, -484\rangle_{A \rightarrow B}$	+269	-274
B^1	$\langle\{r_{14}, r_{15}, r_{16}\}, \{r_2, r_3, r_4\}, +7\rangle_{B \rightarrow A}$	-209	+217
A^2	$\langle\{r_2, r_3, r_4\}, \{r_{14}, r_{15}, r_{16}\}, -245\rangle_{A \rightarrow B}$	+30	-35
B^2	$\langle\{r_{14}, r_{15}, r_{16}\}, \{r_2, r_3, r_4\}, +133\rangle_{B \rightarrow A}$	-82	+90
...
A^6	$\langle\{r_2, r_3, r_4\}, \{r_{14}, r_{15}, r_{16}\}, -217\rangle_{A \rightarrow B}$	+2	-7
B^6	$\langle\{r_{14}, r_{15}, r_{16}\}, \{r_2, r_3, r_4\}, +223\rangle_{B \rightarrow A}$	+8	+0
A^7	<i>accept</i>		

In this example, reframing allows agent *B* to achieve his goal (which he did not achieve in the BO case) while agent *A* achieves her goal for a cost of 1191 (which is cheaper than in the BO case). The next section presents a more systematic study of the qualitative and quantitative advantages of IBN.

6 Simulations and results

The aim of our simulations is to evaluate the eventual benefits of the BR strategy over the BO one. To this end, we conducted experiments using three different scenarios:

1. *The buyer-seller case*: Only one agent has a goal to achieve. Consequently, only this agent can potentially need some resources. He (the buyer) will have to buy these from the other agent (the seller) and a deal consists of exchanging resources for money.
2. *General negotiation with asymmetric reframing*: The two agents are pursuing goals and may want to exchange sets of resources. However, only one of the agents is capable of initiating a reframing.
3. *General negotiation with symmetric reframing*: same as above, but with the two agents being able to initiate a reframing after an unsuccessful bargaining. This is the case described so far in this paper. BR agents are using the meta-strategy number 6 as defined in Sect. 4.3.2.

These three settings are of increasing complexity and, as we will see, some properties of the simpler buyer-seller case no longer hold in more complex settings. The second scenario allows us to investigate the impact of asymmetry on the model's behavior, while the third one covers the general case. Before presenting and discussing the results of these experiments, we first detail our experimental settings and describe the various dimensions that we use to compare BO and BR negotiations.

6.1 Experimental settings

In order to evaluate and characterize the hypothetical benefit(s) of using reframing, we conducted simulations of bilateral negotiation between agents for which errors on the valuation of resources not owned was varied from -70% to $+70\%$ by steps of 5% . The standard deviation of this error was set to zero²⁰ and the potential benefit b was set to 20% . For each combination of errors, 100 different IBN domains were generated and for each of them, BO and BR negotiations were conducted. In other words, each figure showing our results hereafter has been generated by some 168200 negotiations.²¹ Each negotiation consists of a number of instances of the bargaining protocol and in the BR case some instances of the reframing one as well.

The simulations were conducted using randomly generated IBN domains in which plans are trees with a branching factor of 2, of depth 3, and the number of alternative decompositions is 4 (but only for the root decomposition). This results in 4 different plans being generated per agent for a total of 32 resources in the system.

6.2 Qualitative and quantitative dimensions when comparing BO and BR

In order to assess the proposed model, we need to define a metric that can be used to compare negotiation outcomes. When evaluating the benefit of reframing, it is important to differentiate qualitative differences from quantitative ones.

The main *qualitative dimension* of the outcome of a particular negotiation is whether a deal is reached or not. This is related to the main qualitative interest of the agents: that is, achieving

²⁰ For the sake of simplicity, we assumed that the error made by a given agent over all the resources not owned was constant.

²¹ We have 29 possible errors level for both agents and we run 100 simulations in the two conditions, for a total of $29 \times 29 \times 200 = 168200$ negotiations.

their goals. Indeed, while sometimes a deal enables both agents to achieve their goals, at least one agent's goal will be achieved when a deal is reached (e.g., a buyer-seller case). Note that deals are sufficient but not always necessary since agents can sometimes achieve their goals without reaching a deal (e.g., if they have plans without lacking resources, like agent A in the BO case of the example presented in Sect. 5.2).

With each particular choice of simulation parameters being used to instantiate two negotiations (a BO and a BR one), four different combinations of goal achievement are possible for each agent:

1. The agent is not able to reach his goal in the BO case, while he does in the BR case. In this situation, there is a qualitative benefit of using reframing on top of bargaining.
2. The agent reaches his goal in the BO case but not in the BR one. In that case, there is a qualitative loss associated with the use of reframing.
3. The agent reaches his goal in the BO and the BR cases. In this situation, we will be interested in comparing the two cases to see if BR provides any quantitative benefits over BO.
4. The agent does not reach his goal in either the BO or the BR case. While the agent does not reach his goal, he might act as a seller and quantitative benefit may still be worth looking for.

In this study, we are interested in testing whether reframing allows agents to achieve their goals more often (by reaching deals when BO is unsuccessful). This qualitative dimension can be characterized using a measure of the benefits in terms of the number of simulations that fall into case 1 minus the number of simulations that fall into case 2. We call this measure the *benefit in goals*. We will also present the total number of deals reached in the BO case as a baseline on which those benefits in goals sit. This information will be useful to calculate the statistical significance of our results.

Other qualitative dimensions can help to assess the usefulness of reframing. A particular reframing can lead to (a) a proposal of an alternative and cheaper plan. If this information is new, that will entail (b) an update of the valuation function of the agent. We will use the number of occurrences of (a) and (b) to assess the impact of reframing on the agents' preferences (through the manipulation of their valuation functions).

In cases where both the BO case and the BR case lead to the same outcome in terms of a particular agent's goal achievement (cases 3 and 4 above), quantitative dimensions can be used to compare BO and BR strategies. In particular, when an agent achieves its goal in both the BO and BR case (case 3 above), the costs of the plans used to achieve the agents' goals may be different. We denote as *benefit on plan's cost* the difference between the costs of these plans.

Definition 11 (*Benefit on plan's cost*) When plans P_i^1 and P_i^2 allow agent i to achieve his goal in the BO and BR cases respectively, the benefit on plan's cost is defined as: $planbene f_i = cost_i(P_i^2) - cost_i(P_i^1)$

Another interesting quantitative dimension is the subjective benefit gained by an agent during the negotiation. Often, a deal is made before the agents reach their least preferred offer. The offer evaluation function (Definition 5) defines the difference in valuation between the current offer and their least preferred one. In case of acceptance of the offer, this function returns the subjective benefit of the agent. When at least one strategy allows a deal to be reached, the difference between those subjective benefits can be calculated. Note that the no-deal situation entails an evaluation of zero ($eval_i(\emptyset) = 0$).

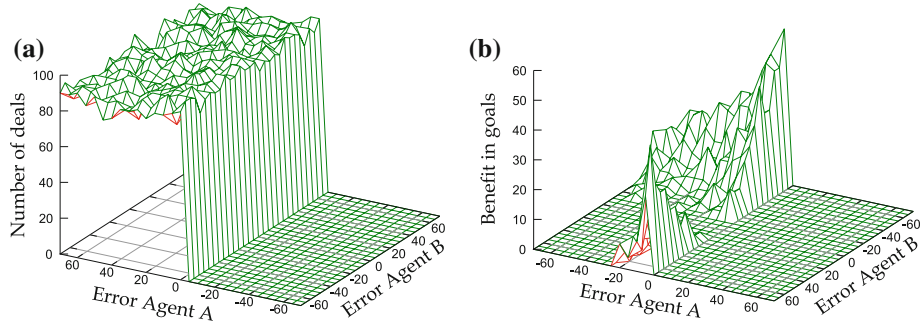


Fig. 4 Qualitative results for agent *A* in buyer-seller scenarios. From left to right, we have: **a** the number of deals reached in the BO case and **b** the *benefit in goals* in case BR is used. Notice that the abscissa and ordinate graduations of the graph **a** have been inverted for better readability when compared to **b** or the other figures

Definition 12 (*Subjective benefit in deals*) Given p^{BO} and p^{BR} , the two outcomes in a BO and BR case respectively (accepted offers or the no-deal \emptyset), the subjective benefit in deals for agent i is defined as: $dealbenef_i = eval_i(p^{BR}) - eval_i(p^{BO})$

6.3 The buyer-seller scenario

The first experiment considers buyer-seller scenarios. Agent *A* is pursuing a goal and is likely to need some of agent *B*'s resources in order to achieve it. Agent *A* will thus act as a buyer. Agent *B* does not have a domain-related goal to accomplish, and will thus never ask for resources. Agent *B* is however interested in selling resources as long as she makes some profit (i.e., subjective benefit in deals, Definition 12). In those scenarios, only agent *B* can initiate a reframing that can, if successful, lead agent *A* to update her valuation function.

Figure 4a shows the results of this experiment in terms of the qualitative dimensions defined in Sect. 6.2. When agent *A* over-evaluates *B*'s resources, a deal will always be reached (left half of the Fig. 4a). Note that deals are reached 100% of the time and the noise that can be observed corresponds to the fact that sometimes agent *A* owns all the resources of her preferred plan so that no negotiation is needed (and no deal is reached).

Whenever *A* over-evaluates *B*'s resources, there will be no difference between the BO and the BR cases since agent *A* always makes an over-evaluated initial proposal that agent *B* accepts directly.²² The first bargaining will always succeed, and reframing will never take place (right half of Fig. 4b).

Conversely, when agent *A* under-evaluates *B*'s resources (right half of Fig. 4a), no deal is reached in the BO case while significant benefits in goals can be made using the BR strategy (left half of Fig. 4b). In those cases, agent *A* only achieves her goal when she has a plan for which she owns all the resources.

For the quantitative results, in cases where both BO and BR allow *A* to achieve her goal, there are substantial benefits on the plan's cost (Fig. 5c) and subjective benefit on deals (Fig. 5d) when reframing is used. Those benefits vary according to agent *B*'s errors about *A*'s resources' valuations. These errors shape agent *B*'s cost preference relation over *A*'s plans. The benefits for *A* are greater when agent *B* underestimates agent *A*'s resources. Fig. 6a

²² Note that, in the buyer-seller scenario, the buyer (i.e., agent *A*) is always the initiator of the bargaining because she is the only agent requesting resources.

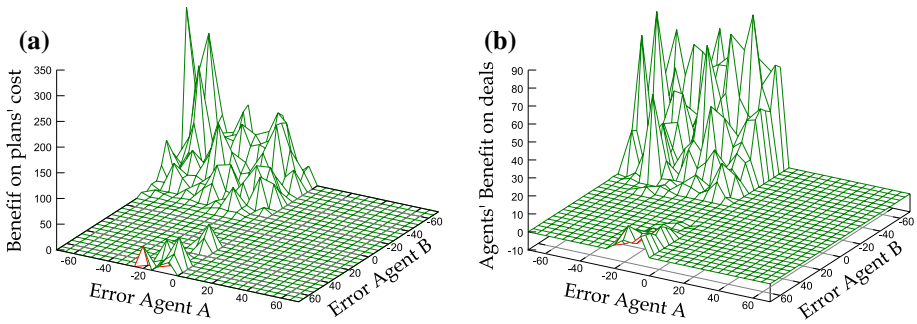


Fig. 5 Quantitative results for agent *A* in buyer-seller scenarios, when both BO and BR allow achieving the goal: **a** the mean of the benefit on the cost of the plan (Definition 11) and **b** the mean difference in subjective benefits from deals (Definition 12)




Fig. 6 Buyer-seller scenario—results for agent *A* (which acts as the buyer): **a** number of valuation updates, **b** number of times only BO allows *A* to achieve her goal

explains this phenomenon by showing that more successful reframing (i.e., more valuation function updates) occurs in that case. This explanation can be refined by having a closer look at the negotiation logs.

When agent *A* underestimates *B*'s resources, she tends to prefer plans which include more missing resources, as these are undervalued. Symmetrically, when *B* overestimates *A*'s resources, he will also tend to prefer the plans involving more of his own resources. In that case, *B* will have fewer alternatives to propose to *A* during the reframing since they tend to agree on the preferable plans (according to their cost preference relations). Consequently, the reframing has less impact on *A*'s valuation update, as shown by Fig. 6a.

Conversely, when *B* underestimates *A*'s resources, the plans involving more of *A*'s resources will start being evaluated as the most preferable ones. In those cases, the agents' preferred plans diverge and the reframing has more impact on the negotiation as shown by the rising number of valuation updates in Fig. 6a.

Furthermore, Fig. 6b shows that there is no case in which BO does better than BR. However, there are clearly cases in which BR does better than BO, as shown in Fig. 4b. As far as *B* is concerned, BR also allows him to improve his benefit, as shown in Fig. 7. These experimental evidences allow us to conclude that in those buyer-seller scenarios, BR Pareto-dominates BO. That is, there is no case in which BO gives a better result than BR, while BR sometimes improves the result of BO. Even if our experimental results give more precise

Fig. 7 Buyer-seller scenario—results for agent *B* (which acts as the seller): difference in benefit made using the BR and the BO strategy

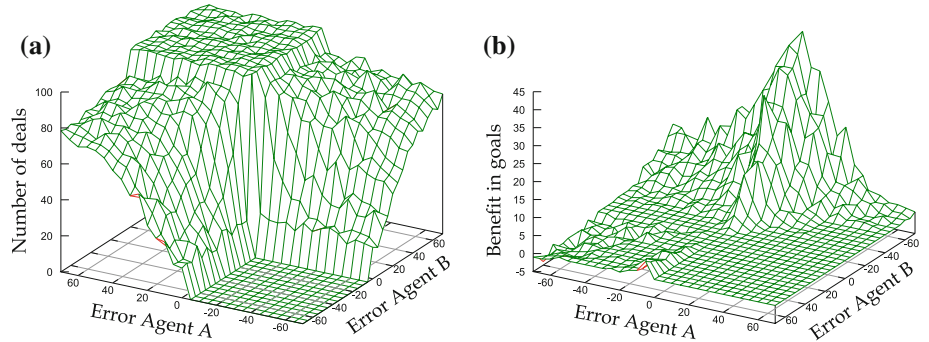



Fig. 8 General case with asymmetric reframing—qualitative results. From *left to right*, we have: **a** the number of deals reached using BO and **b** the number of goals achieved with BR and not with BO for agent *A*. Notice that the abscissa and ordinate graduations of the graph of **a** have been inverted when compared to **b** or the other figures for better readability

indications about the benefit obtained, in this particularly restricted case, this result could have been analytically reached, as noticed in previous work on IBN [38].

Unfortunately, this result will no longer hold in more complex settings. That is what justifies the use of empirical simulation to study the model.

6.4 General case with asymmetric reframing

Our second experiment looks at the case in which both agents have a goal to accomplish, but only agent *B* is allowed to initiate a reframing. This means that agent *A* is the only one who can update her evaluation function (only agent *B* can propose cheaper alternatives). While this case is more complex than the previous one, it simplifies the model by ensuring that agent *B*'s preference relations, \succ_B^c and \succ_B^b , over both his own plans and agent *A*'s plans, will not change. This setting also gives an opportunity to study the consequence of asymmetry in the model.

Figure 8a shows the number of deals reached in the BO case while Fig. 8b and Fig. 10a show the qualitative benefits of using reframing for agents *A* and *B* respectively. Note that, when both agents overestimate the other's resources, both agents tend to give too much for the resources and a deal is always reached during the first bargaining. In that case—as in the similar case in the buyer seller scenarios—the agents will never have a chance to use reframing and there is no qualitative or quantitative difference between BO and BR.

In the asymmetric scenarios, the consequences of reframing are asymmetric as well, as illustrated by Fig. 8b, Figs. 9 and 10. When *B* underestimates *A*'s resources, the reframing is

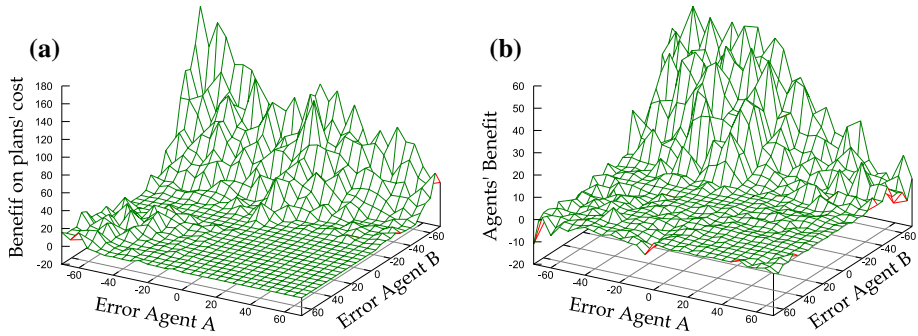


Fig. 9 General case with asymmetric reframing—quantitative results for agent A. From left to right we have: **a** the mean of agent A's benefit on the cost of the plan used and **b** the mean of the difference in subjective benefit made during the deal




Fig. 10 Comparing BO and BR for agent B. From left to right we have: **a** the number of goals achieved with BR and not with BO, and when both BO and BR achieved the goal, **b** the mean of the benefit on the cost of the plan used and **c** the mean of the difference in subjective deal benefit

in most cases beneficial to both agents. However, when agent B overestimates A's resources less than A underestimates B's resources (see the triangular shape, projected on the $x \times y$ plane, in Fig. 10a), advising agent A through the reframing strategy can lead to negative results for B. This can be explained as follows.

When agent A updates her valuation function val_A after a successful reframing, she takes all her plans into consideration again. In this asymmetric scenario, agent B will by assumption never update his valuation function (since A is not conducting any reframing) causing the preferred plan to be dropped after any unsuccessful bargaining. Agent B can thus easily run out of plans, while agent A still has plans to try. This is a bad side-effect of our bargaining agent behavioral model as presented in Sect. 4.3.1. The BO strategy has been designed having in mind a symmetric case and would have to be adapted in order to cope with asymmetric scenarios like those considered here. For example, the agent can stick with his least preferred plan rather than discarding it as is the case with the current strategy. When the agent has no more plans to consider, he becomes a seller for the end of the negotiation, thus renouncing his goal. These negative results, as well as some that are not generated by such asymmetry, are further discussed and exemplified in Sect. 7.1.

Overall, it is no surprise that agent A benefits from these asymmetric scenarios more than B since A is the only one to receive information about alternatives evaluated as cheaper by B. The next section describes our last experiment.

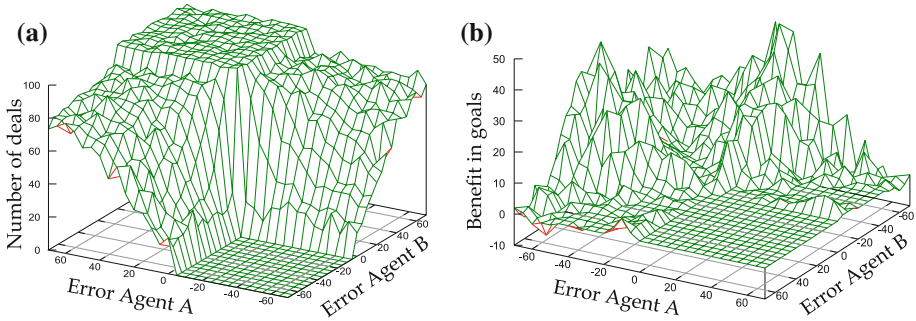


Fig. 11 The general case—qualitative results. From *left to right* we have: **a** deals reached in the BO case and **b** benefit in term of number of goals achieved in the BR case. Notice that the abscissa and ordinate graduations of the graph of **a** have been inverted when compared to **b** or the other figures for better readability

6.5 The general, symmetric case

The general case corresponds to the model described in the rest of this paper. Both agents have a goal to achieve and both can initiate reframing. Meta-strategy number 6 (as presented in Sect. 4.3.2) is used to sequence bargainings and reframings.

6.5.1 Frequency and quality of the deals

The results for the qualitative dimension are summed up by Fig. 11b, which presents the benefit in terms of the difference in the number of goals achieved between BR and BO negotiations. This shows the interest of reframing as a strategy that allows agents to reach deals more often than with bargaining only. Here again, and for the reasons detailed in previous sections, this difference—that is, the qualitative advantage—disappears when the agents overestimate the resources. It also suggests that the advantage of reframing can be made bigger by using meta-strategies other than the current one (number 6 in the list of Sect. 4.3.2). For example, a strategy that would allow reframing before the first bargaining may improve the overall results. However, these other meta-strategies introduce unfairness in certain cases as explained in Sect. 4.3.2. Exploring the various other meta-strategies discussed in Sect. 4.3.2 is left as future work.

Another qualitative dimension is whether a particular reframing is successful (an alternative plan is proposed) and whether it is taken into account by the agent (the information is actually new to the agent: i.e., the reframing is followed by an update of the valuation function of the agent). Fig. 13a shows the mean over the number of updates of the agents' valuation functions per negotiation.

Quantitative dimensions of the quality of deals can also be used to compare BO and BR strategies:

- Subjective benefit made by the agents: Fig. 12d presents the mean of the difference between the subjective benefit made by the agents in deals reached by the BO vs. BR agents;
- Benefit in terms of the cost of the plan(s) enabled by the deal made: Fig. 12c presents the mean of the difference in cost of the plans.

These aspects are correlated with the second qualitative dimension, i.e., successful reframing has a positive effect on the quality of the deal. In conclusion, not only are more deals

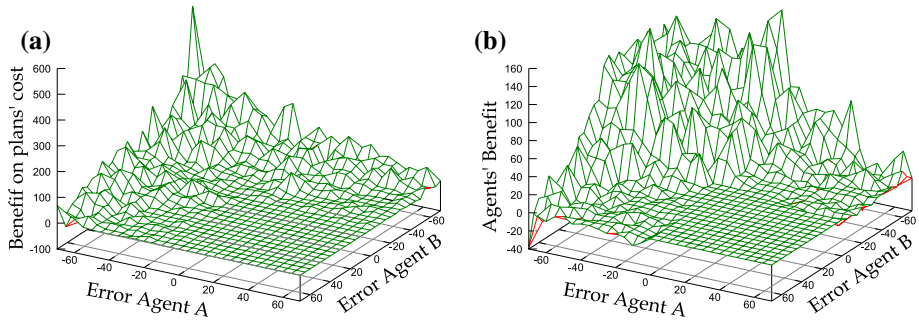


Fig. 12 The general case—quantitative results (sum for both agents). When both agents achieve their goals we have: **a** benefit on the cost of the plan used, **b** difference in benefit of the agents between the two conditions




Fig. 13 **a** Mean of the number of valuation updates per negotiation, **b** mean over the number of messages per negotiation in the BO case and **c** mean over the number of messages exchanged in the BR case

reached in the BR case but also the deals reached are of better quality for the agents. Notice that we present both agents' results together and the figures are more or less symmetrical, which indicates that the proposed negotiation process is not biased in favor of one of the parties.

6.5.2 Negotiation complexity

In order to measure the overhead of using reframing, we assumed that the size of messages is bounded by a constant and we measured the number of messages used in BO negotiations (Fig. 13b) and in BR negotiations (Fig. 13c). These numbers are correlated with the number of bargainings²³ made in the first case (bounded to four with the domain values used for the simulations) and the number of bargainings and reframings made in the second one (Fig. 13a gives a lower bound for both in the BR cases).

The cost of reframing in terms of communication is clear. However, the bottleneck of the system is the number of alternative plans as calculated by Algorithm 1 presented in Sect. 3 which is exponential in the number of alternatives per goal. This is multiplying the number of possible negotiations and affects both the BO and BR negotiation systems. This result regarding the complexity of the search space is quite common and usual solutions and heuristics to circumvent it would apply [51].

²³ This denotes the number of times the bargaining protocol is used and should not be confused with the number of bargaining rounds within an instance of the bargaining protocol.

6.6 Synthesis of the results

Table 1 synthesizes and summarizes our results. For each of the three scenarios considered in our study (buyer seller, asymmetric and symmetric), the results for a given type of error made on the estimations by each agent are presented. Each result corresponds to the average of the difference between the BO and BR conditions in each experiment, and the standard deviation is indicated in parenthesis. More precisely, the table presents our results using the following columns:

- **Case:** This column numbers the lines for ease of reference.
- **Situation:** For each line, this column indicates which of the three situations is considered. The three possibilities are: (a) the buyer-seller scenario (only *A* may need *B*'s resources, i.e., *A* is the buyer, *B* the seller), (b) the general case with asymmetric reframing (only *B* can initiate a reframing, only *B* can thus recommend cheaper alternatives and only *A* can update her valuation function) and (c) the general symmetrical case.
- **Error A:** For each line, this column indicates what type of errors agent *A* is doing in her original evaluation of the resources she does not own. The following types of errors can be considered:
 1. *Underestimate:* This encompasses all the cases in which the agent somehow underestimates the value of the resources not owned: that is, when the error is between 0 and -70% of the valuation maintained by the owner.
 2. *Overestimate:* This encompasses all the cases in which the agent overestimates the value of the resources not owned: that is, when the error lies between 0 and $+70\%$ of the valuation maintained by the owner.
 3. *Both* (noted $-$): This encompasses all the cases, regardless of the error made (between -70 and $+70\%$).
- **Error B:** idem for agent *B*.
- **Agent:** For each line, this column indicates which agent is considered.
- **goals:** This column indicates the average benefit in goals for the agent considered (standard deviation is indicated in parenthesis). These qualitative results can be read as a percentage indicating how often BR allows the agent to reach his goal where BO did not.
- **bcost:** This column indicates the average benefit in terms of the plan cost for the agent considered as introduced in Definition 11 (standard deviation is indicated in parenthesis). This quantitative advantage can occur only when both BR and BO were successful.
- **bdeals:** This column indicates the average benefit in terms of the resources exchanged as defined in Definition 12 for the agent considered (standard deviation is indicated in parenthesis). This quantitative advantage is computed as the difference in all negotiations (0 is used when no deal is reached).
- **p-value:** This column presents an upper boundary on the *p*-value resulting from the statistical significance test as described below.

While 27 cases, or lines,²⁴ could have been shown, only 19 relevant ones have been listed to avoid redundant information.

The table also indicates the statistical significance of the results presented in this paper and by doing so backs up all the general claims made, while presenting our detailed analysis of these results in previous sections. Since for each experiment, the BO and BR conditions have been conducted under the very same settings, domain parameters and with the same

²⁴ Three situations, three types of errors for both agents, that is: $3 \times 3 \times 3 = 27$ settings.

agents, single-tailed (or one-sided)²⁵ paired *t*-tests have been conducted. The *t*-test compares two means (given their standard deviation) and tests a null hypothesis stating that they do not differ. The *p*-value indicates the probability of obtaining results at least as extreme as the ones that were actually observed, assuming that the null hypothesis is true. The lower the *p*-value, the higher the chances are that the claim that one condition (BR in our case) is a significant improvement over the other (BO in our case) is true. In our case, all the results presented are highly statistically significant (because they have very low *p*-values, 0.01 being the standard in inferential statistics). This is a direct consequence of the number of experiments conducted. Indeed, each result is an average of either:

1. 42050 negotiations, that is 21025 comparisons between BO and BR conditions. Note that this corresponds to an average over one quadrant in the Figure presenting the results.
2. 84100 negotiations, that is 42050 comparisons between BO and BR conditions. Note that this corresponds to an average over half a Figure presenting the results.
3. 168200 negotiations, that is 84100 comparisons between BO and BR conditions. This corresponds to an average over a whole Figure.

These large numbers allow us to draw conclusion without ambiguity.

Lines 1 to 8 of Table 1 generalise and synthesise results for the buyer seller case, presented in Figs. 4, 5 and 7 and can be read as follows:

1. Lines 1 and 2: when agent *A* overestimates the value of the resources not owned, there is neither qualitative nor quantitative difference between BO and BR for agent *A*. There is no difference in benefits between BO and BR for agent *B* (bgoals and bcost are irrelevant in that case since *B* is the seller).
2. Lines 3 and 4: when both agents underestimate the value of the resources not owned, there are significant qualitative and quantitative advantages for *A* and a significant quantitative advantage for *B* in terms of benefits made when a deal is reached.
3. Lines 5 and 6: when agent *A* underestimates the value of the resources not owned, there are significant qualitative and quantitative advantages for *A* and a significant quantitative advantage for *B* in terms of benefits made when a deal is reached.
4. Lines 7 and 8: overall (that is, independently of the errors made by the agents knowing that these would be comprised between -70% and $+70\%$), there are significant qualitative and quantitative advantages for *A* and a significant quantitative advantage for *B* in terms of benefits made when a deal is reached.

Lines 9 to 15 generalise and synthesise results for the general asymmetric case, presented in Figs. 8, 9, and 10 and can be read as follows:

1. Lines 9 and 10: when both agents overestimate the value of the resources not owned, there is neither qualitative nor quantitative difference between BO and BR for either agent.
2. Lines 11 and 12: when *B* underestimate the value of the resources not owned, there are significant qualitative or quantitative advantages for both agents. The asymmetric reframing possibility favors *A* for the qualitative and plan cost dimension (and the difference is significant).
3. Lines 13 and 14: when agent *A* underestimates the value of the resources not owned and *B* overestimates these, there are significant qualitative and quantitative disadvantages for agent *B*. This negative result is further analyzed and explained in Sects. 6.4 and 7.1.

²⁵ A two-tailed test is used when one just wants to conclude that there is a difference, while a one-sided test draw conclusion about the sign of that difference.

Table 1 Synthesis of the results

Case	Situation	Error A	Error B	Agent	bggoals	bcost	bdeals	p-value
1	Buyer seller	Overestimate	-	A	0 (0)	0 (0)	0 (0)	1
2	Buyer seller	Overestimate	-	B	-	-	0 (0)	1
3	Buyer seller	Underestimate	Underestimate	A	13.4 (7.3)	56.4 (45.8)	47.1 (32.1)	<0.000001
4	Buyer seller	Underestimate	Underestimate	B	-	-	39.7 (18.6)	<0.000001
5	Buyer seller	Underestimate	-	A	4.3 (3.9)	26.4 (28.2)	24.3 (29.1)	<0.000001
6	Buyer seller	Underestimate	-	B	-	-	20.2 (21.3)	<0.000001
7	Buyer seller	-	-	A	2.2 (2.6)	13.2 (23.4)	12.2 (17.3)	<0.000001
8	Buyer seller	-	-	B	-	-	10.1 (14.3)	<0.000001
9	Asymmetric	Overestimate	Overestimate	A	0 (0)	0 (0)	0 (0)	1
10	Asymmetric	Overestimate	Overestimate	B	0 (0)	0 (0)	0 (0)	1
11	Asymmetric	-	Underestimate	A	8.4 (6.2)	38.4 (25.8)	21.7 (16.1)	<0.000001
12	Asymmetric	-	Underestimate	B	1.8 (2.1)	8.9 (3.8)	22.3 (17.1)	<0.000001
13	Asymmetric	Underestimate	Overestimate	B	-1.8 (1.7)	-3.2 (3.0)	-0.1 (0.1)	<0.000001
14	Asymmetric	-	-	A	4.2 (4.3)	19.7 (16.8)	10.9 (10.1)	<0.000001
15	Asymmetric	-	-	B	0.2 (1.9)	3.7 (4.1)	10.9 (9.1)	<0.000001
16	Symmetric	Overestimate	Overestimate	Both (sum)	0 (0)	0 (0)	0 (0)	1
17	Symmetric	Underestimate	-	Both (sum)	11.1 (8.5)	87.4 (65.8)	21.7 (16.1)	<0.000001
18	Symmetric	-	Underestimate	Both (sum)	10.8 (8.7)	88.1 (64.3)	22.0 (15.9)	<0.000001
19	Symmetric	-	-	Both (sum)	8.1 (8.3)	67.6 (65.4)	16.1 (16.0)	<0.000001

Numbers in parenthesis are standard deviations

4. Lines 14 and 15: overall (that is, independent of the errors made by the agents knowing that these are comprised between -70% and $+70\%$), there are significant qualitative and quantitative advantages for both agents. The qualitative and plan cost advantage is clearly less important for B .

Lines 16 to 19 generalise and synthesise results for the general case, presented in Figs. 11 and 12 and can be read as follows:

1. Line 16: when both agents overestimate the value of the resources not owned, there is neither qualitative nor quantitative difference between BO and BR for either agent.
2. Lines 17 and 18: when at least one of the agents underestimates the value of the resources not owned, there are significant qualitative or quantitative advantages for both agents.
3. Line 19: overall (that is, independent of the errors made by the agents knowing that these are comprised between -70% and $+70\%$), there are significant qualitative and quantitative advantages to use BR for both agents.

While many other claims can be made from our experimental results, these are the main ones we wanted to emphasize to clarify our contribution.

One last but important remark is that while the quantitative results are somehow tied to the specifics of the proposed bargaining protocol and strategies, the qualitative results are not. Indeed, these results emphasize the fact that reframing allow deals to be reached in the BR cases where no deal is possible in the BO cases. These qualitative results hold independently of the bargaining protocol and strategies used. Indeed, when no deal is possible, this is because the bargaining interval—i.e., the intersection of the preference intervals of the agents—is empty and no bargaining protocol will reach a deal in that case. Our bargaining protocol and strategy guarantee that a deal will be found whenever the bargaining interval is not empty. This property should hold for any reasonable bargaining procedure. The emphasis of the paper is thus not on finding the best bargaining protocol or strategy but rather on making a point about the benefit of reframing in general.

7 Discussion

7.1 Negative results

The shapes of the figures presenting our experimental results highlight the non-linearity and the complexity of the proposed model. This shows the inherent complexity of interactions between agents with partial, imperfect, imprecise or erroneous knowledge. While the BR case Pareto-dominates the BO case in the buyer-seller scenario (as shown in Sect. 6.3), this is not the case in the other, more general, cases. In fact, there are a few cases in which the figures indicate negative results, i.e. combinations for which BO outperforms BR.

Indeed, there are cases in which BO agents succeed in reaching their goals while BR agents fail (for the same set of domain data). For example, suppose both the agents have three plans available for reaching their goal and the preference ordering for their own plans is defined as follows: $P_A^1 \succ_A^c P_A^2 \succ_A^c P_A^3$ for A and $P_B^1 \succ_B^c P_B^2 \succ_B^c P_B^3$ for B . In the BO case, the following series of bargainings will occur where we assume that the third one is a success:

A's selected plan	B's selected plan	Bargaining
P_A^1	P_B^1	<i>failure</i>
P_A^2	P_B^2	<i>failure</i>
P_A^3	P_B^3	<i>success</i>

In the BR case, the first bargaining of this sequence will still fail (like in the BO case). The agents will then try to reframe. Suppose the cost preference ordering of agent *B* over the plans for agent *A* is: $P_A^3 \succ_B^c P_A^1 \succ_B^c P_A^2$. In this case, agent *B* will suggest that plan P_A^3 is actually cheaper than plan P_A^1 . Agent *A* will update her valuation function and take all the plans into account again while agent *B* will not reconsider plan P_B^1 . The rest of the sequence of bargainings will then be:

A's selected plan	B's selected plan	Bargaining
P_A^3	P_B^2	<i>failure</i>
P_A^1	P_B^3	<i>failure</i>
P_A^2	\emptyset	<i>failure</i>

A successful reframing will generally change the combination of bargainings that occur. This can lead to a situation in which no bargaining will succeed, resulting in a failure of the agents' goals.

Note that these cases (as well as related ones) are rare and that this effect is counter-balanced by the fact that the successful reframing itself increases the likelihood of reaching an agreement. These cases are also linked to the chosen meta-strategy as discussed in previous sections. Meta-strategies that would preclude reframing to perform worse or sub-optimal in some cases by revisiting plans in a more systematic fashion may have other consequences, e.g., higher communication overhead, longer negotiations or even infinite loops. Inquiring solution to fix this problem is left as future work.

7.2 Loops

There are also cases in which the proposed (recursive) reframing method loops. Of course, as such loops can be detected, the agents have been programmed to stop trying to reframe and go back to a BO strategy in such cases.

Suppose that the preference ordering for agent *A* is: $P_A^1 \succ_A^c P_A^2 \succ_A^c P_A^3$ and the preference ordering agent *B* has over the plans of *A* are: (1) cost preference ordering: $P_A^2 \succ_B^c P_A^3 \succ_B^c P_A^1$ and (2) profit preference ordering: $P_A^3 \succ_B^b P_A^2 \succ_B^b P_A^1$. We assume that agent *A* has no suggestions for agent *B* and that all the bargainings fail. In the following table, we can see the sequence of agent *A*'s cost preference relation resulting from the iteration of the updating of the valuations of the resources' values according to *B*'s recommendations.

Preferences of agent <i>A</i> (subjective costs)
$P_A^1(400) \succ^c P_A^2(500) \succ^c P_A^3(550)$
$P_A^3(475) \succ^c P_A^1(476) \succ^c P_A^2(500)$
$P_A^1(476) \succ^c P_A^2(487.5) \succ^c P_A^3(488.5)$
$P_A^3(481) \succ^c P_A^1(483) \succ^c P_A^2(487.5)$
...

The costs are converging and then start looping as the plans get reordered in a circular way. Several extensions of the model are possible to handle these (rare) cases in a more elegant way, for example by enabling agents to memorize and reason about the coherence of the other agent's recommendations. Different update functions will also have various impacts on this phenomenon.

These two phenomena are a direct consequence of making the agents' preferences dynamic in the context of uncertain and possibly erroneous valuation knowledge. Moving away from the specifics of the previous discussions next section summarizes and generalizes our results.

7.3 Stability

Stability refers to the fact that agent should not have an incentive to deviate from the agreed-upon strategies proposed by the given negotiation framework [58]. We do not want agent designers (e.g., companies) to have an incentive to build their agents with different, manipulative, strategies.

In analytical frameworks, like game-theoretic ones, simplification are made, usually though assumptions of complete and perfect knowledge, so that the optimal behavior can be computed (eventually off-line). When that is the case, the optimal behavior is publicly revealed and there is nothing better to do than just carrying it out. In case of incomplete information, it is sometimes possible to design mechanisms in the sense that it is always individually rational for each rational to behave non strategically (in particular truthfully). Such mechanisms are said to be incentive compatible.

On one hand, given our assumptions and the realism of our model—including the presence of erroneous information—we did not succeed in showing stability of our approach. As in any other automated negotiation framework, there is indeed a number of manipulations that are possible during the bargaining (e.g., decoy bargaining, premature bargaining failure, premature bargaining acceptance, delayed bargaining acceptance) or the reframing (e.g., hidden goal, phantom goal, wrong alternative, false information, cheap talk) [47].

On the other hand, we did not find any case in which an agent has a clear, locally computable, incentive to deviate from the prescribed strategies using some of the above mentioned manipulations. The same complexity that precludes us to prove properties of stability precludes the agents to develop rational deviation strategies.

Further inquiring the classical game theoretic properties of negotiation frameworks—individual rationality, symmetry, efficiency and stability—is left as future work.

8 Conclusion and future work

The traditional form of automated negotiation, characterized by position-based approaches, is restricted to bargaining. This consists of an exchange of offers by the agents trying to accommodate each other's preferences until a deal is acceptable to both parties or the negotiation terminates unsuccessfully. These approaches tend to view the object of the negotiation and the agents' preferences as being fixed. By focusing on interests to be satisfied rather than positions to be won, reframing—a particular IBN strategy—allows the agents to search the space of negotiation objects (rather than the space of deals for a particular exchange of items).

Previous works have claimed that IBN improves the likelihood and the quality of agreements. Our main aim was to test that hypothesis in the context of a formal, computational

model. Despite a number of assumptions made explicit throughout the text, the presented model is very general and captures both the usual buyer-seller negotiations as well as more general cases of exchange of sets of resources with monetary compensation. The notion of consumable resources has been left unspecified to further enhance this generality and the applicability of the approach.

In conclusion, the results presented in this paper show that reframing significantly improves the quantity of successful negotiations (i.e., negotiations that allow the agents to achieve their goals). Furthermore, when the negotiation restricted to bargaining is already successful, then the use of the reframing strategy tends to reduce the cost of the plans and improves the benefit made during the deal. It is however crucial to notice that these qualitative and quantitative advantages are not regular results of an analytical nature, but statistical results that hold in general rather than in every case.

In his survey of models for automated negotiation, Buttner [4] notes that among the 74 articles reviewed; only 7 take advantage of argumentation-based approaches, only 12 consider some sort of incomplete information about the negotiated item(s), and none does both. More generally, the model of automated negotiation proposed in this paper partially addresses the four limitations attached to position-based approaches and mentioned in Sect. 2.1.3. Indeed, in the proposed model:

1. The agents do not have any a priori knowledge (not even stochastic) of the other's utility function or goal(s) (no restriction to complete, common information).
2. The agents have imperfect (i.e., erroneous) information on the value of the resources not owned. The error made by the agents on the evaluation of the resources not owned has been varied to explore this dimension (no restriction to perfect information).
3. The reframing strategy takes advantage of the communication and cognitive capabilities of goal-driven artificial agents (no restriction to an exchange of offers).
4. The reframing strategy entails that the agents' preferences are updated dynamically. Both the cost and the benefit preference relations may be affected by reframing dialogues (no restriction to static preferences).

The expressivity and realism gained, when added to the robustness to incomplete and erroneous information about the object of the negotiation and the opponent, gives a promising ground for building real-world applications in the context of electronic commerce. This model can be seen as a first step towards filling the gap between formal and computational approaches for IBN (and ABN) [3].

In giving a first empirical evaluation of reframing, this work builds foundations for future work. More experiments will be done to extend the current results by comparing different strategies for updating the agents evaluation of the resources they do not own as well as different ways to combine reframing and bargaining. Experiments will look at different types of domain structure as well. Other bargaining and reframing protocols will be developed allowing agents to exchange information about know-how in order to relax assumptions 6 and 5 respectively.

The model will also be extended to encompass positive and negative interactions between goals, thus allowing the representation of more domains and cases. Such a study has already been started in a simplified variant (i.e., one that can be analytically addressed) of that model [42]. To progress toward real-world applications of these exiting results, further studies, experiments and validations are also under investigation in the field of human-computer negotiation (see [13] for preliminary results).

Acknowledgments We would like to thank the anonymous reviewers for their helpful comments and the Center for High Performance Computing of The University of Melbourne for their services and support. This research was made possible by the Australian Research Council and the Natural Sciences and Engineering Research Council of Canada.

References

1. Barbuceanu, M., & Lo, W.-K. (2000). A multi-attribute utility theoretic negotiation architecture for electronic commerce. In *Proceedings of the fourth international conference on autonomous agents* (pp. 239–246).
2. Bierman, H. S., & Fernandez, L. (2005). *Game theory with economic applications*. Prentice Hall: Addison-Wesley.
3. Buttner, R. (2006). A classification structure for automated negotiations. In *Proceedings of the 4th international workshop on multi-agent systems in e-business (MASEB'06)* (pp. 523–530). IEEE Computer Society.
4. Buttner, R. (2006). The state of the art in automated negotiation models of the behavior and information perspective. *International Transactions on Systems Science and Applications (ITSSA)*, 1(4), 351–356.
5. Cohen, S. P. (2002). *Negotiating skills for managers*. New York: McGraw-Hill.
6. Conitzer, V., & Sandholm, T. W. (2004). Computational criticisms of the revelation principle. In *Proceedings of the 5th ACM conference on electronic commerce (EC-04)* (pp. 262–263).
7. Dastani, M., de Boer, F., Dignum, F., & Meyer, J.-J. (2003). Programming agent deliberation: An approach illustrated using the 3APL language. In *Proceedings of the second conference on autonomous agents and multi-agent systems (AAMAS'03)* (pp. 97–104).
8. Dastani, M., van Riemsdijk, M. B., & Meyer, J.-J. (2005). Programming multi-agent systems in 3APL. In *Multi-agent programming: Languages, platforms and applications* (pp. 10–45). Heidelberg: Springer-Verlag.
9. Endriss, U. (2006). Monotonic concession protocols for multilateral negotiation. In *Proceedings of the 5th international joint conference on autonomous agents and multiagent systems (AAMAS'06)* (pp. 392–399). New York: ACM Press.
10. Faratin, P., Sierra, C., & Jennings, N. R. (2002). Using similarity criteria to make trade-offs in automated negotiations. *Artificial Intelligence*, 142(2), 205–237.
11. Fisher, R., & Ury, W. (1983). *Getting to yes: Negotiating agreement without giving in*. New York: Penguin Books.
12. Fudenberg, D., Levine, D., & Tirole, J. (1985). Infinite horizon models of bargaining with one sided incomplete information. In *Game theoretic models of bargaining*. Cambridge: Cambridge University Press.
13. Gal, Y., D'souza, S., Pasquier, P., Rahwan, I., & Abdallah, S. (2009). The effects of goal revelation on computer-mediated negotiation. In *Proceedings of the annual meeting of the cognitive science society (CogSci)* (pp. 2614–2619). The Netherlands: Amsterdam.
14. Gal, Y., & Pfeffer, A. (2007). Modeling reciprocity in human bilateral negotiation. In *National conference on artificial intelligence (AAAI)*. British Columbia: Vancouver.
15. Governatori, G., Dumas, M., Hofstede, A., & Oaks, P. (2001). A formal approach to protocols and strategies for (legal) negotiation. In *Proceedings of the 8th international conference on artificial intelligence and law* (pp. 168–177). New York: ACM Press.
16. Grosz, B., Kraus, S., Talman, S., & Stossel, B. (2004). The influence of social dependencies on decision-making: Initial investigation with a new game. In *Proceedings of the third international joint conference on autonomous agents and multi-agent systems (AAMAS'04)* (pp. 782–789).
17. Harsanyi, J. C., & Selten, R. (1972). A generalized Nash solution for two-person bargaining games with incomplete information. *Management Science*, 18(5), 80–106.
18. Heiskanen, P., Ehtamo, H., & Hamalain, R. P. (2001). Constraint proposal method for computing pareto solutions in multi-party negotiations. *European Journal of Operational Research*, 133(1), 44–61.
19. Hiltrop, J. M., & Udall, S. (1995). *The essence of negotiation*. Englewood Cliffs, NJ: Prentice Hall.
20. Jennings, N. R., Faratin, P., Lomuscio, A. R., Parson, S., Sierra, C., & Wooldridge, M. (2001). Automated negotiation: Prospects, methods, and challenges. *Journal of Group Decision and Negotiation*, 2(10), 199–215.
21. Jonker, C. M., Robu, V., & Treur, J. (2007). An agent architecture for multi-attribute negotiation using incomplete preference information. *Autonomous Agents and Multi-Agent Systems*, 15(2), 221–252.

22. Karunatillake, N. C., Jennings, N. R., Rahwan, I., & McBurney, P. (2009). Dialogue games that agents play within a society. *Artificial Intelligence*, 173(9-10), 935–981.
23. Karunatillake, N. C., Jennings, N. R., Rahwan, I., & Ramchurn, S. D. (2006). Managing social influences through argumentation-based negotiation. In *Proceedings of the fifth international joint conference on autonomous agents and multiagent systems (AAMAS'06)* (pp. 426–428). New York: ACM Press.
24. Kowalczyk, R. & Bui, V. (2001). On constraint-based reasoning in e-negotiation agents. In Dignum, F. & Cortés, U. (Eds.) *Agent-mediated electronic commerce III*. Lecture Notes in Computer Science, (Vol. 2003, pp. 31–46). Berlin: Springer Verlag.
25. Kraus, S. (2001). *Strategic negotiation in multi-agent environments*. Cambridge MA, USA: MIT Press.
26. Kraus, S., & Sycara, K. (1998). Reaching agreement through argumentation: A logical model and implementation. *Artificial Intelligence*, 104(1–2), 1–69.
27. Larson, K., & Sandholm, T. W. (2001). Bargaining with limited computation: Deliberation equilibrium. *Artificial Intelligence*, 132, 183–217.
28. Lesser, V., Decker, K., Wagner, T., Carver, N., Garvey, B., & Horling, B., et al. (2004). Evolution of the GPGP/TAEMS domain-independent coordination framework. *Autonomous Agents and Multi-Agent Systems*, 9(1), 87–143.
29. Lin, R., Kraus, S., Wilkenfeld, J., & Barry, J. (2006). An automated agent for bilateral negotiation with bounded rational agents with incomplete information. In *Proceeding of the 17th European conference on artificial intelligence (ECAI)* (pp. 270–274).
30. Lin, R., Kraus, S., Wilkenfeld, J., & Barry, J. (2008). Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence Journal*, 172(6–7), 823–851.
31. Linhart, P. B., Radner, R., & Satterthwaite, M. A. (1991). *Bargaining with incomplete information*. San Diego: Academic Press.
32. Luo, X., Jennings, N. R., Shadbolt, N., Leung, H. F., & Ho-man Lee, J. (2003). A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments. *Artificial Intelligence*, 148, 53–102.
33. Matos, N. & Sierra, C. (1998). Evolutionary computing and negotiating agents. In *Agent mediated electronic commerce*. Lecture Notes on Artificial Intelligence, (Vol. 1571, pp. 126–150). Heidelberg: Springer-Verlag.
34. McKersie, R. B., Sharpe, T., Kochan, T., Eaton, A., Strauss, G., & Morgenstem, M. (2008). Bargaining theory meets interest-based negotiations: A case study. *Industrial Relations: A Journal of Economy and Society*, 47(1), 66–96.
35. Nash, J. F. (1950). The bargaining problem. *Econometrica*, 18, 155–192.
36. Osborne, M. J., & Rubinstein, A. (1994). *A course in game theory*. Cambridge: The MIT Press.
37. Parson, S., Sierra, C., & Jennings, N. R. (1998). Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3), 261–292.
38. Pasquier, P., Dignum, F., Rahwan, I., & Sonenberg, L. (2006). Interest-based negotiation as an extension of monotonic bargaining in 3APL. In *Proceedings of the ninth Pacific rim international workshop on multi-agents (PRIMA 2006)*. Lecture Notes in Artificial Intelligence (LNAI) (Vol. 4088, pp. 569–580). Guilin, China: Springer-Verlag.
39. Pasquier, P., Hollands, R., Dignum, F., Rahwan, I., & Sonenberg, L. (2007). An empirical study of interest-based negotiation. In *Proceedings of the ninth international conference on electronic commerce (ICEC)* (pp. 339–348). New York: ACM Press.
40. Pruitt, D. G. (1981). *Negotiation behavior*. New York: Academic Press.
41. Rahwan, I. (2004). *Interest-based Negotiation in Multi-Agent Systems*. PhD thesis, Department of Information Systems, University of Melbourne, Melbourne, Australia.
42. Rahwan, I., Pasquier, P., Sonenberg, L., & Dignum, F. (2009). Formal analysis of interest-based negotiation. *Annals of Mathematics and Artificial Intelligence*, 55(3–4), 253–276.
43. Rahwan, I., Pasquier, P., Sonenberg, S., & Dignum, F. (2007). On the benefits of exploiting underlying goals in argumentation-based negotiation. In *Proceedings of the 22nd conference on artificial intelligence (AAAI)*. Vancouver, Canada: AAAI Press.
44. Rahwan, I., Ramchurn, S., Jennings, N., McBurney, P., Parsons, S., & Sonenberg, L. (2003). Argumentation based negotiation. *Knowledge Engineering Review*, 18(4), 343–375.
45. Rahwan, I., Sonenberg, L., & Dignum, F. (2003) Towards interest-based negotiation. In *Proceedings of the 2nd international joint conference on autonomous agents and multiagent systems (AAMAS'03)* (pp. 773–780). New York: ACM Press.
46. Raiffa, H. (1982). *The art and science of negotiation*.: Harvard University Press.

47. Rosenschein, J. S., & Zlotkin, G. (1994). *Rules of encounter: Designing conventions for automated negotiation among computers*. Cambridge, Massachusetts: MIT Press.
48. Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50(1), 97–109.
49. Rubinstein, A. (1985). A bargaining model with incomplete information about time preferences. *Econometrica*, 53, 1151–1172.
50. Rubinstein, A. (1997). *Modeling bounded rationality*. Cambridge MA, USA: MIT Press.
51. Russell, S. & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. Series in Artificial Intelligence, 2nd (ed.). Englewood Cliffs: Prentice Hall.
52. Sandholm, T. (1999). Distributed rational decision making. In G. Weiss (Ed.), *Multiagent systems: A modern approach to distributed artificial intelligence* (pp. 201–258). Cambridge, MA: The MIT Press.
53. Sandholm, T., & Vulkan, N. (1999). Bargaining with deadlines. In Hendler, J. & Subramanian, D. (Eds.), *Proceedings of the 16th national conference on artificial intelligence (AAAI-1999), Menlo Park CA, USA*. (pp. 44–51).
54. Wolfstetter, E. (1996). Auctions: An introduction. *Journal of Economic Surveys*, 10, 367–420.
55. Wooldridge, M. (2001). *An Introduction to multiagent systems*. New York: Wiley.
56. Young, P. (1991). *Negotiation analysis*. Ann Arbor: University of Michigan Press.
57. Zeng, D., & Sycara, K. (1997). Benefits of learning in negotiation. In *Proceedings of the 14th national conference on artificial intelligence and 9th innovative applications of artificial intelligence conference (AAAI-97/IAAI-97), Menlo Park, USA* (pp. 36–42, 27–31). Vancouver, Canada: AAAI Press.
58. Zlotkin, G. & Rosenschein, J. S. (1989). Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of the eleventh international joint conference on artificial intelligence (IJCAI-89)* (pp. 912–917). Detroit: MI.