

P. Pasquier, R. Hollands, I. Rahwan, F. Dignum, L. Sonenberg. Exploiting Hierarchical Goals in Bilateral Automated Negotiation: Empirical Study. Proceedings of the AAMAS International Workshop on Agent Mediated Electronic Commerce (AMEC), Honolulu, Hawaii, USA.

Exploiting Hierarchical Goals in Bilateral Automated Negotiation: Empirical Study

Philippe Pasquier¹, Ramon Hollands³, Frank Dignum³, Iyad Rahwan^{2,4}, and
Liz Sonenberg¹

¹ University of Melbourne, Australia

² British University of Dubai, UAE

³ Utrecht University, The Netherlands

⁴ (Fellow) University of Edinburgh, UK

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.

1 Introduction and Motivation

Approaches to automated negotiation can be classified in three categories [3]: (1) game theoretic (2) heuristic and (3) argumentation based. 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) and have erroneous estimations of the value of the resources not owned. This implies that we aim to produce good rather than optimal solutions. On the other hand, heuristic approaches suffer from two main limitations [3]: (1) the only feedback given to unacceptable proposals is either a counter-proposal or a rejection and (2) the preferences of the agents are statically defined which entails that the set of issues under negotiation is static. Our approach proposes a model of recursive *reframing* where constructive feedbacks are given according to the underlying goals revealed by the agents. This is a type of argumentation-based (ABN) strategy implementing interest-based negotiation (IBN) that can change the agents preferences and thus the issues under negotiation.

While ABN and IBN have been the focus of many publications, very few (if any) empirical evaluations have been provided yet [5]. IBN advocates the idea that parties can increase the likelihood and quality of an agreement by exchanging information about their underlying goals and alternative ways to achieve them that influence the agents' preferences [4]. This work advances the state of the art in automated negotiation by testing this hypothesis empirically.

In order to reach that aim, we define a negotiation model suited to agents with hierarchical goals (Section 2). This model includes a bargaining protocol and monotonic concession strategy (Section 3.1), a recursive reframing protocol

and strategy (Section 3.2) as well as the metastrategy to articulate those two strategies in the agent behavior (Section 3.3). 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 (Section 5).

2 Agents with hierarchical goals

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

- $\mathcal{A} = \{A_1, \dots, A_n\}$ is the set of agents;
- $\mathcal{G} = \{G_1, \dots, G_m\}$ is the set of all possible goals;
- $goal_i \in \mathcal{G}$ is the root goal of agent i ,⁵
- $Res = \{r_1, \dots, r_p\}$ is the set of resources;
- $res_i \subseteq Res$ is the set of resources owned by agent i ;
- $sub : \mathcal{G} \times 2^{\mathcal{G} \cup Res}$ is the relationship between goals and their decomposition in sub-goals and resources needed to achieve them,⁶
- $val_i : Res \rightarrow \mathbb{R}$ is a function that returns, for each agent, its valuation of its own resources, as well as estimated value of resources not owned by itself.

In that context, agents maintain preference intervals over the selling values of the resources they own and about the acquisition values of the resources they do not possess.

Definition 2 (Preferences).

- $prefmin_i : Res \rightarrow \mathbb{R}$ is a function that returns the less 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 .

Note that $prefmin$ gives what is called in economics [1] 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 as a buyer of the resource). Agents being rational, it follows that:

- $\forall i \in \mathcal{A}, \forall x \in res_i, prefmax_i(x) \geq prefmin_i(x)$
- $\forall i \in \mathcal{A}, \forall x \notin res_i, prefmax_i(x) \leq prefmin_i(x)$

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 $prefmin$ and $prefmax$.

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:

$$\forall i \in \mathcal{A}, \forall x \in Res, prefmin_i(x) = val_i(x)$$

⁵ Throughout the paper, we assume that each agent has a single goal. Multiple goals can be expressed by a single goal that has a single alternative decomposition.

⁶ This is a relation, not a function, to allow expressing that a goal may be fulfilled by multiple sets of *alternative* sub-goals/resources

Other choices are possible. This one means: (1) that agents may be willing to give up profit on the resources they own but are not ready to do or give anything for less than its subjective value and symmetrically, (2) they will not give more than their valuation for the resources they do not own. In other words, the agents are ready to concede their gain up to their reservation price and this one is equal to their subjective evaluation.

We also assume a fixed margin of potential benefit for each agent expressed as a percentage.

Assumption 2 (Potential benefit)

The so-call potential benefit of the an 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 he is a buyer). Another set of assumptions constrains the structure of 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);
- $\bigcup_{i \in \mathcal{A}} \text{res}_i = \text{Res}$ (all the resources are owned);
- Money is part of the resources and it is the only divisible one. In the present approach, money will be treated like a quantity and will be expressed using real numbers.⁷

In this paper, we restrict the negotiation to two agents that have complementary resources. 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. We consider the following assumptions:

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 owned is its value. Note that - a posteriori - the cost of using a resource is given by its value if it is owned and by its acquisition cost if it is a negotiated resource.

Assumption 5 (Shared vs. private knowledge)

We assume that all agents have 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 , 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 others’ goals.

In that context, agents can generate all the possible plans to achieve a particular goal.

Definition 3 (Plan and cost). *A plan P_n^i generated by agent i for achieving a goal G_0 is a tree such that:*

⁷ This is not contradicting the previous statements since the money owned by an agent is not the money owned by another one.

- 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:
 $\text{sub}(G, \{x_1, \dots, x_m\})$;
- Each leaf node x is a resource: $x \in \text{Res}$

We note $\text{needed}(P_n^i)$ the set of leaf nodes of the plan P_n^i and $\text{missing}_i(P_n^j)$ the subset of $\text{needed}(P_n^j)$ not owned by agent i ($\text{missing}_i(P_n^j) = \text{needed}(P_n^j) \setminus \text{res}_i$). More generally, we note missing_i the set of resources that an agent i is willing to acquire.

The estimated cost of a plan P_n^i for agent i is:⁸

$$\text{cost}_i(P_n^j) = \sum_{x \in \text{needed}(P_n^j)} \text{val}_i(x)$$

The potential benefit that an agent i can make on a plan P_n^j - in selling the resources he own that are needed for the plan - is defined by:

$$\text{benef}_i(P_n^j) = \sum_{x \in \text{needed}(P_n^j) \cap \text{res}_i} \text{prefmax}_i(x) - \text{prefmin}_i(x)$$

We note \succ_i^c and \succ_i^b the preference ordering of the plans according to these estimated costs and benefits.

The use of goals (or super-goals) with subgoals gives a great expressivity to the model [6]. Since, we do not over-specify this choice, related representations using the same type of arborescent structures (like tasks, sub-tasks and super-tasks) can still be captured.

Assumption 6 (Goal and plan) We assume that no overlap exists between agents' needed resources: $\forall n, \text{needed}(P_n^i) \cap \text{needed}(P_n^j) = \emptyset$, or between their plans' root or sub-goals: $\forall n, m, \forall x \in P_n^i \Rightarrow x \notin P_m^j$

This last assumption is realistic in many domains where agents operate in separate but complementary sub-domains. Relaxing this assumption would raise the problem of positive and negative interaction between goals and conflicts for resources. We plan to address these issues as future work.

3 The Negotiation Framework

In order to enable agents to use both bargaining and reframing, one needs to define appropriate communication protocols. Figure 1 thus 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.

3.1 Bargaining: Protocol and Strategy

Part (a) of Figure 1, describes the proposed bargaining protocol. In order to ensure that at least the initiator agent needs a negotiation dialogue to occur, we assume that $\text{missing}_i \neq \emptyset$. The bargaining protocol initiated by agent i with agent j is divided in two parts as follows:

⁸ Note that we assume that the resources will be consumed by the plan execution as specified by assumption 4.




Fig. 1. UML 2.0 specification of the bargaining and reframing protocols.

- Part one (negotiating the resources to be exchanged): Each agent discloses the set of resources that he wants⁹:
 $inform_i(need(missing_i))$ and $inform_j(need(missing_j))$
- Part two (bargaining over the payment):
 1. i makes a first offer (Definition 7)
 2. 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 (see Definition 6) is met in which case bargaining terminates with a deal;
 - *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 that 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 the eventual (and subjectively evaluated by i) difference in value.

Definition 5 (Proposal evaluation).

⁹ Assumptions 6 simplify this part of the protocol, relaxing them (which is future work), will complexify it.

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

$$\begin{aligned} \text{eval}_i(p_i) &= \sum_{x \in S_{\text{wanted}}} \text{val}_i(x) - \sum_{x \in S_{\text{given}}} \text{val}_i(x) - \text{Payment} \\ \text{eval}_j(p_i) &= \text{Payment} + \sum_{x \in S_{\text{given}}} \text{val}_j(x) - \sum_{x \in S_{\text{wanted}}} \text{val}_j(x) \end{aligned}$$

Definition 6 (Acceptance of a proposal). An agent j will accept a proposal $p_i = \langle S_{\text{wanted}}, S_{\text{given}}, \text{Payment} \rangle_{i \rightarrow j}$ iff $\text{eval}_j(p_j) - \text{eval}_j(p_i) \leq \rho$, where p_j is the next proposal to be issued by j and ρ is a strictly positive real standing for the deviation tolerated by j ;¹⁰

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

$$p_j^1 = \langle \text{missing}_j, \text{missing}_i, \text{Payment} \rangle_{j \rightarrow i}$$

, where Payment is defined as being:

$$\sum_{j, r_i \in \text{missing}_j} \text{prefmax}_j(r_j) - \sum_{j, r_j \in \text{missing}_i} \text{prefmax}_i(r_j)$$

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

Definition 8 (Concession strategy). Given a proposal $p_i = \langle S_{\text{wanted}}, S_{\text{given}}, \text{Payment} \rangle_{i \rightarrow j}$ received by j consequently to his previous proposal p_j^t , the next proposal to be issued by j is $p_j^{t+1} = \langle S'_{\text{wanted}}, S'_{\text{given}}, \text{Payment}' \rangle_{j \rightarrow i}$ where:

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

$$\text{eval}_j(p_j^{t+1}) = \frac{(\text{eval}_j(p_j^t) + \text{eval}_j(p_i))}{2}$$

If $\text{eval}_j(p_j^{t+1}) \geq 0$ and $\text{money}_j > \text{Payment}'$ then:

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

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 the agents are in that situation, the bargaining ends without reaching 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} .

¹⁰ ρ is usually quite small and is here just to avoid infinite length bargaining. It can be expressed as a percentage of $\text{eval}_j(p_j)$, typically between 0.1% and 5% depending on the domain.

3.2 Reframing: Protocol and Strategy

Part (b) of Figure 1, describes the 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 $missing_j$. The agent j can then (1) disclose his set of underlying goals $g_j = \{x_1, \dots, x_n\}$ motivating the will to acquire $missing_j$ while clarifying the sub-plan(s) selected for achieving it or (2) rejects 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 $\succ_j i^c$) to be 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 in a rational way according to \succ_i^b , that is the one with the highest potential benefit is disclosed first.¹¹

On the reception of the information that there is an alternative plan P_j^2 for achieving a goal from g_j that 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 are fixed. 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 a 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.
- 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).¹²

Other choices corresponding to strategies of type (1) or (2) are worth studying but are left as future work. In any case, this is a cooperative answer since it makes i 's statement true in j 's model which assumes that j trusts i .

¹¹ In case of equal potential benefit, a random choice is made.

¹² In the absence or more information, it is assumed that the alternative plan is just a little bit cheaper (where “a little bit” is represented by ρ).

3.3 Agents 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.

Bargaining only agents (BO) A BO agent's execution cycle is as follows:

1. The agent generates the 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 Section 3.1;
4. If the bargaining fails, she withdraws the current plan and proceeds with the next cheapest plan.

The process terminates when there is no plan left.¹³

Reframing capable agents (BR) BR agents' execution cycle extends the one of BO agents with the reframing capabilities described in Section 3.2. 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 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 function is updated (or not) according to the last reframing, the following meta-strategies give different outcomes:

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

In this paper, BR agents will use the meta-strategy number 6. The agents will start with a (complete) bargaining. If bargaining fails, both agents will attempt¹⁴ reframing before initiating a new bargaining. All reframings have to terminate before the agents (eventually) update their valuation functions and a new bargaining is to be initiated. The rationale for this choice is that such "parallel" reframing does not give any 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 to future work.

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

¹³ Since each agents has a finite number of plans and the monotonic concession protocol is known to terminate [2], negotiations between BO agents always terminates.

¹⁴ Following a buyer-seller bargaining only one agent can attempt a reframing.

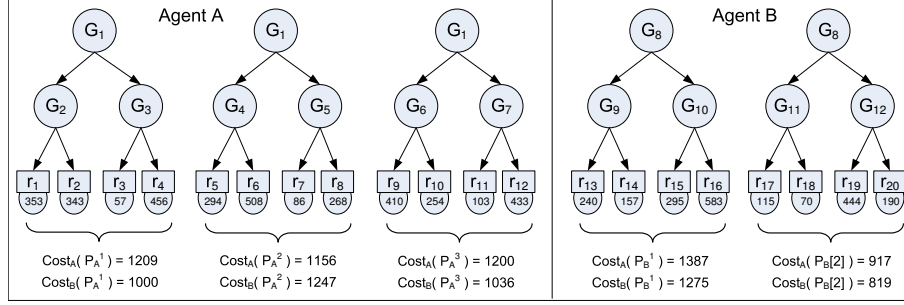


Fig. 2. Plans initially generated by agents *A* and *B*.

4 Simulation and Example

4.1 Parameters of experimentation

Parameters about the domain It is clear that the structure of the domain will influence the results obtained in the experiments envisioned to compare BO and BR negotiations. In particular, the complexity and richness of the domain will have impact on the usefulness of using reframing (reframing being completely useless when no alternatives are available whatsoever).

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

Parameters related to the agents For each agent introduced in the systems, 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 doing 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 *var*) are expressed as percentages. Finally - for both agents - the preferences boundaries (*prefmin* and *prefmax*) are calculated for each agent 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).

4.2 Detailed Example

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

- $\mathcal{A} = \{A, B\}$ is a set of agents;
- $\mathcal{G} = \{G_1, \dots, G_{12}\}$ is a set of all possible goals;
- $goal_A = G_1$ and $goal_B = G_8$;
- $Res = r_1, \dots, r_{20}$;
- $sub = \{(G_1, \{G_2, G_3\}), \dots \text{ as given by Figure 2 } \dots (G_{12}, \{r_{11}, r_{12}\})\}$;
- $res_A = \{r_1, r_7, r_8, r_9, r_{10}, r_{11}, r_{12}, r_{14}, r_{15}, r_{16}, r_{17}, r_{18}, r_{19}, r_{20}, money = 2000\}$;
- $res_B = \{r_2, r_3, r_4, r_5, r_6, r_{13}, money = 1500\}$;
- The values of $val_A, val_B, prefmin_A, prefmin_B, prefmax_A$ and $prefmax_B$, has been generated with $error_A = error_B = 0\%$ and $var_A = var_B = 70\%$, with a potential benefit of 20%:
 $[353, 388]_A^{r_1}, [308, 343]_A^{r_2}, [51, 57]_A^{r_3}, [411, 456]_A^{r_4}, [265, 294]_A^{r_5}, [457, 508]_A^{r_6},$
 $[86, 94]_A^{r_7}, [268, 295]_A^{r_8}, [410, 451]_A^{r_9}, [254, 278]_A^{r_{10}}, [103, 113]_A^{r_{11}}, [433, 476]_A^{r_{12}},$
 $[220, 244]_A^{r_{13}}, [371, 408]_A^{r_{14}}, [200, 220]_A^{r_{15}}, [468, 515]_A^{r_{16}}, [114, 126]_A^{r_{17}}, [95, 105]_A^{r_{18}},$
 $[470, 517]_A^{r_{19}}, [154, 169]_A^{r_{20}}$ and $[47, 174]_B^{r_1}, [314, 346]_B^{r_2}, [66, 73]_B^{r_3}, [445, 490]_B^{r_4},$
 $[432, 475]_B^{r_5}, [427, 470]_B^{r_6}, [80, 89]_B^{r_7}, [222, 247]_B^{r_8}, [262, 328]_B^{r_9}, [171, 214]_B^{r_{10}},$
 $[72, 90]_B^{r_{11}}, [323, 404]_B^{r_{12}}, [240, 264]_B^{r_{13}}, [126, 157]_B^{r_{14}}, [266, 295]_B^{r_{15}}, [524, 583]_B^{r_{16}},$
 $[103, 115]_B^{r_{17}}, [63, 70]_B^{r_{18}}, [400, 444]_B^{r_{19}}, [171, 190]_B^{r_{20}}$

In this example, we assume that $\rho = 1$. The agents generate all the possible plans to achieve their main goals along with their costs as shown in Figure 2. 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 being P_A^2 and P_B^2 respectively (see Figure 2), we have: $missing_A = \{r_5, r_6\}$ 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 A^1 and B^1 (Definition 7). Then, according to the concession strategy (Definition 8), A and B repeat themselves 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 resources, agent A will achieve his goal on his 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^3	$\langle \{r_{14}, r_{15}, r_{16}\}, \{\}, +1029.5 \rangle_{B \rightarrow A}$	-9.5	+5.5
A^5	<i>reject</i>		

In this case, only agent A succeeds in achieving his 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,




Fig. 3. From left to right we have the: (a) the number of goals achieved with BR and not with BO, when both BO and BR achieved the goal, we have: (b) the mean of the benefit on the cost of the plan used and (c) the mean of the difference in subjective benefit of the agents.

they both compute the possible plans for the other agents to achieve them along with their costs 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^5 , see Figure 2 for the costs). Note that B proposes P_A^1 rather than P_A^3 according to \succ_B^b (P_A^3 does not allow any benefit).

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

Agent A will update his 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 $[265, 332]_{A^2}^{r_2}$, $[37, 64]_{A^3}^{r_3}$, $[356, 445]_{A^4}^{r_4}$), and the cost of P_A^2 becomes 1192. Agent A 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	accept		

In this example, reframing allows agent B to achieve his goals while agent A achieves his goal for a cost of 1191 (which is a cheaper than in the BO case).




Fig. 4. (a) Mean of the number of valuation update per negotiation, (b) mean over the number of message per negotiation in the BO case and (c) mean over the number of messages exchanged in the BR case.

5 Experimental Results

In order to evaluate and characterize the hypothetical benefit of using reframing, we have been conducting 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 null and the potential benefit b was set to 20% . For each combination, 100 BO negotiations have been conducted and 100 BR ones (assorted in couples with the same initial domain and valuation values). In other words, each curve showing our results hereafter has been generated by some $29 * 29 * 200 = 168200$ negotiations. Each negotiation is made 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 with a branching factor of 2, trees of depth 3 and number of alternative 4 (but only for the root decomposition), that is 4 different plans generated per agent for a total of 32 resources in the system.

5.1 Frequency and Quality of the deals

When comparing BO and BR negotiations to evaluate the benefit of reframing, it is important to differentiate qualitative differences from quantitative ones.

The main *qualitative dimension* of a negotiation is whether a deal is reached or not. This is related to the main qualitative interest of the agents, that is achieving their goals. Indeed, while sometimes a deal allows both agents to achieve their goals, at least one agent will achieve his goal in case a deal is reached (a buyer-seller case). Note that deals are sufficient but not always necessary since agents can achieve their goals without reaching a deal if they have some plan without missing resources. In our case, it is interesting to see if reframing allows agents to achieve their goals more often (by reaching deals with BR when BO is unsuccessful)

The results for this dimension are summed up by Figure 3.a that shows 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 allowing the agents to reach deals more often than with bargaining only. This difference - that is the qualitative advantage - disappears when the agents overestimate the resources. Indeed in that case (first quarter in the Figures), the first bargaining

tends to succeed in both cases (and no reframing occurs). While this improves the Figures readability, it also suggests that the advantage of reframing can be made bigger by using other meta-strategies than the current one (number 6 in the list of Section 3.3). Exploring the various meta-strategies discussed in Section 3.3 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). Figure 4.a shows the means over the number of updates of the agents valuation functions per negotiation.

In the case where both BO and BR allow an agent to achieve his goal, *quantitative dimensions* of the quality of deals can be used to compare BO and BR strategies:

- Benefit in terms of the cost of the plan(s) enabled by the deal made: Figure 3.b presents the mean of the difference in cost of the plans;
- Subjective benefit made by the agents: Figure 3.c presents the mean of the difference between the subjective benefit made by the agents in deals reached by the BO vs. BR agents;

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 more deals are reached in the BR case but also the deals reached are of better quality for the agents.

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 (Figure 4.b) and in BR negotiations (Figure 4.c). These numbers are correlated with the number of bargainings 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 (Figure 4.a 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 that is clearly exponential in the number of alternatives per goal and this is common to the BO and BR negotiation systems. This result about the complexity of the search space is not new nor surprising and usual solutions to circumvent it would apply.

The shapes of Figures 3, a, b and c highlight the non-linearity and the complexity of the proposed model. It would be at least hard (if ever possible) to give an analytical account of the advantages of BR over BO. This is due to the inherent complexity of interactions between agents with partial and imperfect knowledge.

6 Conclusion and Future Work

The traditional form of negotiation, characterized by the assertion of opposing positions by the parties (i.e. bargaining), is referred to as position-based negotiation. This tends 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 allows the agents to search the space of negotiation objects (rather than the space of deals for a particular exchange of items). This IBN strategy takes advantage of the communication and cognitive capabilities of goal-driven artificial agents which are ignored by traditional formal approaches to automated negotiation, usually grounded in applied mathematics or micro-economics.

Even though, there is more work required to generalize them, the results presented in this paper allow to conclude that reframing significantly improves the quantity of successful negotiations (i.e. negotiations that allow the agents to achieve their goal). Furthermore, when the negotiation is already successful in the BO case, then the BR 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 analytical nature but statistical results that hold in general rather than in every cases.

In giving a first empirical evaluation of reframing - a particular IBN strategy - this work builds foundations for further experiments. More experiments will be done 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.

Acknowledgments. We would like to thank the High Performance Computing Center of the University of Melbourne and the Australian Research Council for their support.

References

1. H.S. Bierman and L. Fernandez. *Game Theory with Economic Applications*. Addison-Wesley Publishing Compagny, 2005.
2. U. Endriss. Monotonic concession protocols for multilateral negotiation. In P. Stone and G. Weiss, editors, *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006)*, pages 392–399. ACM Press, May 2006.
3. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parson, C. Sierra, and M. Wooldridge. Automated negotiation: Prospects, methods, and challenges. *Journal of Group Decision and Negotiation*, 2(10):199–215, 2001.
4. P. Pasquier, F. Dignum, I. Rahwan, and L. Sonenberg. 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)*, volume 4088 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 569–580, Guilin, China, 2006. Springer-Verlag.
5. I. Rahwan, S. Ramchurn, N. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation based negotiation. *Knowledge Engineering Review*, 18(4):343–375, 2003.
6. I. Rahwan, L. Sonenberg, and F. Dignum. Towards interest-based negotiation. In J. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2003)*, pages 773–780. ACM Press, 2003.