

An Empirical Study of Interest-based Negotiation

Philippe Pasquier
Liz Sonenberg
Dept. of Information Systems
University of Melbourne
Parkville, VIC 3010, Australia
pasquier,
l.sonenberg@unimelb.edu.au

Iyad Rahwan
Institute of Informatics
British University in Dubai
POBox 502216, Dubai, UAE
(Fellow) School of Informatics,
University of Edinburgh,
Edinburgh, UK
irahwan@acm.org

Frank Dignum
Ramon Hollands
Dept. of Information &
Computing Sciences
Utrecht University
Utrecht, The Netherlands
dignum,
rholland@cs.uu.nl

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.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Intelligent Agent, Multiagent Systems*

General Terms

Theory, Algorithms, Experimentation

Keywords

Automated Negotiation, Interest Based Negotiation

1. INTRODUCTION AND MOTIVATIONS

Negotiation is the search for agreement on the exchange (or allocation) of scarce resources among (self-)interested parties. Approaches to automated negotiation have been classified in three categories [4]: (1) game theoretic (2) heuristic and (3) argumentation based.

While interesting conceptually, game theoretic approaches are restricted by strong assumptions which are the price to pay for enabling formal proofs to be made. More specifically in automated negotiation models for electronic commerce, game theoretic approaches generally assume:

- 1 *Common knowledge of the agents utility function*, while “typically, the parties do not know each other’s utility functions with any degree of accuracy. . . Usually, they don’t know each other BATNA [Best Alternative to a Negotiated Agreement]” [10, p.5].
 - 2 *Common (and accurate) knowledge of the valuation of the negotiated object(s)* (common prior), while “ordinary experience seems to indicate that what makes horse races is variation among prior” [6].
- While allowing progress on those two restrictions of game theoretical approaches, heuristic approaches suffer from two other main limitations that they share with game theoretic ones [4]:
- 3 *Agent communication and cognitive capabilities are underused*: The only feedback given to unacceptable proposals is either a counter-proposal or a rejection;
 - 4 *The preferences of the agents are statically defined*: A static preference function is given. It entails that the set of issues under negotiation is static. As noted in [4], the negotiation space does not change dynamically.

In this paper, we propose a new model that overcomes those four limitations.

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 (no common prior).

On the other hand, our approach distinguishes itself from heuristic approaches by proposing a model of recursive *reframing* where constructive feedbacks are given according to the underlying goals revealed by the agents. More precisely, reframing is a sub-type of interest based negotiation strategy that allows the agents to ask for the goal(s) underlying the will for the negotiated item(s) and allow them to give constructive feedbacks (propose alternative plan(s) or reveal information about alternative plans). These elaborated feedbacks can change the agents preferences and thus the issues under negotiation. This strategy may thus allow the agents to reach a deal in this new negotiation space. This strategy is implementing interest-based negotiation (IBN) and can be considered a type of a type of argumentation-based (ABN) strategy resting on deliberative arguments (where the support is made of goals and beliefs and the conclusion is the goal/will to get the negotiated items).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICEC’07, August 19–22, 2007, Minneapolis, Minnesota, USA.
Copyright 2007 ACM 978-1-59593-700-1/07/0008 ...\$5.00.

While ABN and IBN have been the focus of many publications, very few (if any) empirical evaluations have been provided yet [8]. 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 [7]. This work advances the state of the art in automated negotiation by testing this hypothesis empirically.

In order to reach these aims, we define a negotiation model suited to agents with hierarchical goals (Sections 2 and 3). 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 (Sections 4 and 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 ;¹
- $Res = \{r_1, \dots, r_p\}$ 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;²
- $val_i : Res \rightarrow \mathbb{R}$ 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:

¹Throughout 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.

²This is a relation, not a function, to allow expressing that a goal may be fulfilled by multiple sets of *alternative* sub-goals/resources

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

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:

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

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:

- $\forall x \in res_i, prefmax_i(x) = prefmin_i(x) + b * prefmin_i(x)$
- $\forall x \notin res_i, prefmax_i(x) = prefmin_i(x) - b * prefmin_i(x)$

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 the previously described IBN domains.

ASSUMPTION 3 (DISTRIBUTION OF THE RESOURCES).

For distribution of the resources, we assume that:

- $\bigcap_{i \in \mathcal{A}} res_i = \emptyset$ (the resources are not shared);
- $\bigcup_{i \in \mathcal{A}} res_i = Res$ (all the resources are owned);
- *In the present approach, Money will be treated like a quantity and will be expressed using real numbers. A finite amount of Money is part of the resources and it is the only divisible one.*³

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.

³This is not contradicting the previous statements since the money owned by an agent is not the money owned by another one.

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 the cheapest plan}
 $result := plan[i]$ {return the cheapest plan}

Procedure $generate(node, plan[x])$ **Local variables:**

$i := 1$ {number of variations}
 $next$: $node$ {variable of type node}
 $curplan := plan[x]$ {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

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.

2.1 Generating Hierarchical Plans

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

DEFINITION 3 (PLAN AND COST). A 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 note $needed(P_i^n)$ the set of leaf nodes of the plan P_i^n and we note $missing_i(P_j^n)$ 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 note $missing_i$ the set of resources that an agent i is willing to acquire.

The estimated cost of a plan P_i^n for agent i is:⁴

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

The potential benefit that an agent i can make on a plan P_j^n - in selling the resources he owns that are needed for 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 \succ_i^c and \succ_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 [9]. 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 [5]) can still be captured.

ASSUMPTION 6 (GOAL AND PLAN). We assume that no overlap exists between agents’ needed resources, nor between their plans’ root or sub-goals:

- $\forall n, needed(P_i^n) \cap needed(P_j^n) = \emptyset$
- $\forall n, m, \forall x \in P_i^n \Rightarrow x \notin P_j^m$

This last assumption is realistic in many domains where agents operate in separate but complementary sub-domains. 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.

Despite these assumptions, this model is much more general than usual game theoretic approaches. Beyond the absence of knowledge about the opponent utility function, it addresses the case in which agents build their own (i.e. subjective) evaluations of resources according to their own calculus which can be based on: their (eventually uncertain, imprecise or erroneous) beliefs, biased by their interests as well as their past experience, information that has been communicated to them by other agents and so on. . .

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

Those assumptions are thus realistic modelling choices, that correspond to real world situations as usually modelled in micro-economics. A recent synthesis on automated negotiation [2] indicates the need for models that deals with uncertainty about the negotiation object rather than uncertainty about the partner utility function (or type) as it is usually the case.

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




Figure 1: UML 2.0 specification of the bargaining and reframing protocols.

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. A detailed example of negotiation using this framework is presented in Section 4.2.

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 $missing_i \neq \emptyset$. The bargaining protocol initiated by agent i with agent j is divided in two parts as follows:

- Part one: negotiating the resources to be exchanged
Each agent discloses the set of resources that he wants⁵: $inform_i(need(missing_i))$, $inform_j(need(missing_j))$
- Part two: bargaining over the payment

⁵ Assumptions 6 simplify this part of the protocol, relaxing them (which is future work), will complexify it.

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). 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). 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 standing for the deviation tolerated by j ;⁶

DEFINITION 7 (BARGAINING INITIAL PROPOSALS). For each 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 being:

$$\sum_{j, r_i \in missing_j} p_{refmax_j}(r_j) - \sum_{j, r_j \in missing_i} p_{refmax_i}(r_j)$$

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

⁶ 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 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).⁹

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 .¹⁰

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.

3.3.1 Bargaining only agents

A BO agent's execution cycle is represented by Figures 2 and 3 in which the box for reframing would be bypassed and 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 or when a deal is reached.¹¹

3.3.2 Reframing capable agents

Bargaining and reframing (BR) capable agents' execution cycle extends the one of BO agents with the reframing capabilities described in Section 3.2 and Figure 4.

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

⁹In 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).

¹⁰This assumption is to be relaxed. However, the assumption rest on the intuition that it is to the advantage of both agents to be sincere and trust each others statements.

¹¹Since each agents has a finite number of plans and the monotonic concession protocol is known to terminate [3], negotiations between BO agents always terminates.




Figure 4: Execution cycle for the reframing strategy.

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 attempt¹² 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 in our simulations is chosen randomly.

4. SIMULATION AND EXAMPLE

4.1 Parameters of experimentation

4.1.1 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 synthetic domain generation module which allows to define 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 alternatives number 2, we get 8 possible plans to achieve the agent's main goal.

4.1.2 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 doing in evaluating the resources he does not own.

¹²Following a buyer-seller bargaining only one agent can attempt a reframing.

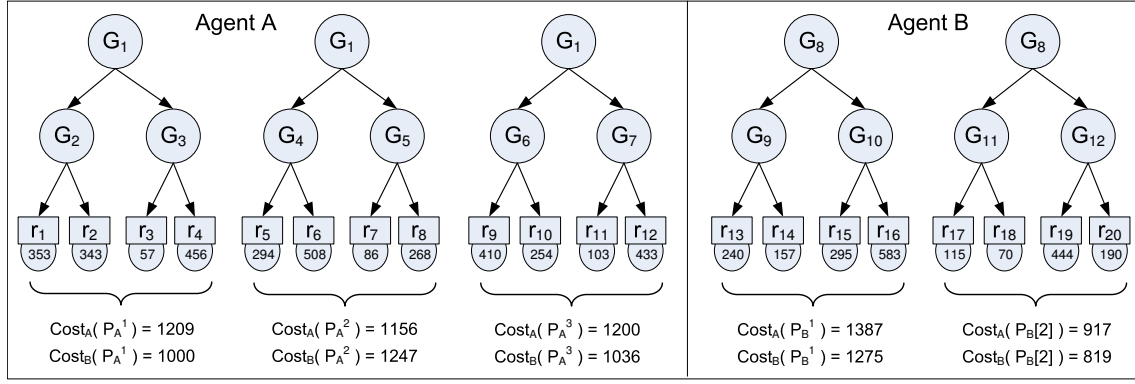


Figure 5: Plans initially generated by agents A and B .

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 $\text{prefmin}_i(r_j) = x$ 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 will 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:

- $A = \{A, B\}$ is a set of agents;
- $\mathcal{G} = \{G_1, \dots, G_{17}\}$ is a set of all possible goals;
- $\text{goal}_A = G_1$ and $\text{goal}_B = G_8$;
- $\text{Res} = r_1, \dots, r_{20}$;
- $\text{sub} = \{(G_1, \{G_2, G_3\}), (G_1, \{G_4, G_5\}), (G_1, \{G_6, G_7\}), (G_1, \{G_8, G_9\}), (G_2, \{r_1, r_2\}), (G_3, \{r_3, r_4\}), (G_4, \{r_5, r_6\}), (G_5, \{r_7, r_8\}), (G_6, \{r_9, r_{10}\}), (G_7, \{r_{11}, r_{12}\}), (G_8, \{G_9, G_{10}\}), (G_8, \{G_{11}, G_{12}\}), (G_9, \{r_{13}, r_{14}\}), (G_{10}, \{r_{15}, r_{16}\}), (G_{11}, \{r_{17}, r_{18}\}), (G_{12}, \{r_{19}, r_{20}\}), (G_{13}, \{r_1, r_2\}), (G_{13}, \{r_3, r_4\}), (G_{14}, \{r_5, r_6\}), (G_{14}, \{r_7, r_8\}), (G_{15}, \{r_9, r_{10}\}), (G_{15}, \{r_{11}, r_{12}\}), (G_{15}, \{G_2, G_3\}), (G_{16}, \{G_4, G_5\}), (G_{16}, \{G_6, G_7\}), (G_{16}, \{r_{13}, r_{14}\}), (G_{16}, \{r_{15}, r_{17}\}), (G_{16}, \{r_{18}, r_{19}\}), (G_{16}, \{G_8, G_9\}), (G_{17}, \{r_1, r_2\}), (G_{17}, \{r_3, r_4\}), (G_{17}, \{r_5, r_6\}), (G_{17}, \{r_7, r_8\}), (G_{17}, \{r_9, r_{10}\})\}$;
- $\text{res}_A = \{r_1, r_7, r_8, r_9, r_{10}, r_{11}, r_{12}, r_{14}, r_{15}, r_{16}, r_{17}, r_{18}, r_{19}, r_{20}, \text{money} = 2000\}$;
- $\text{res}_B = \{r_2, r_3, r_4, r_5, r_6, r_{13}, \text{Money} = 1500\}$;
- The values of val_A , val_B , prefmin_A , prefmin_B , prefmax_A and prefmax_B , have been generated with $\text{error}_A = \text{error}_B = 0\%$ and $\text{var}_A = \text{var}_B = 70\%$ and with a potential benefit of 10%:
 $[353, 388]_A^{r_1}$, $[308, 343]_A^{r_2}$, $[51, 57]_A^{r_3}$, $[411, 456]_A^{r_4}$,
 $[265, 294]_A^{r_5}$, $[457, 508]_A^{r_6}$, $[86, 94]_A^{r_7}$, $[268, 295]_A^{r_8}$,
 $[410, 451]_A^{r_9}$, $[254, 278]_A^{r_{10}}$, $[103, 113]_A^{r_{11}}$, $[433, 476]_A^{r_{12}}$,
 $[220, 244]_A^{r_{13}}$, $[371, 408]_A^{r_{14}}$, $[200, 220]_A^{r_{15}}$, $[468, 515]_A^{r_{16}}$,
 $[114, 126]_A^{r_{17}}$, $[95, 105]_A^{r_{18}}$, $[470, 517]_A^{r_{19}}$, $[154, 169]_A^{r_{20}}$
and $[147, 174]_B^{r_1}$, $[314, 346]_B^{r_2}$, $[66, 73]_B^{r_3}$, $[445, 490]_B^{r_4}$,
 $[432, 475]_B^{r_5}$, $[427, 470]_B^{r_6}$, $[80, 89]_B^{r_7}$, $[222, 247]_B^{r_8}$,
 $[262, 328]_B^{r_9}$, $[171, 214]_B^{r_{10}}$, $[72, 90]_B^{r_{11}}$, $[323, 404]_B^{r_{12}}$,

$$[240, 264]_B^{r_{13}}, [126, 157]_B^{r_{14}}, [266, 295]_B^{r_{15}}, [524, 583]_B^{r_{16}}, \\ [103, 115]_B^{r_{17}}, [63, 70]_B^{r_{18}}, [400, 444]_B^{r_{19}}, [171, 190]_B^{r_{20}}$$

In this example, we assume that $\rho = 1$. The agents generate all the possible plans to achieve their main goals along with their costs as shown in Figure 5. 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 5), we have: $\text{missing}_A = \{r_5, r_6\}$ and $\text{missing}_B = \{r_{17}, r_{18}, r_{19}, r_{20}\}$. 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, 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, 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 5 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 $[265, 332]_A^{r_2}$, $[37, 64]_A^{r_3}$, $[356, 445]_A^{r_4}$), 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	$\{\{r_2, r_3, r_4\}, \{r_{14}, r_{15}, r_{16}\}, -484\}_{A \rightarrow B}$	+269	-274
B^1	$\{\{r_{14}, r_{15}, r_{16}\}, \{r_2, r_3, r_4\}, +7\}_{B \rightarrow A}$	-209	+217
A^2	$\{\{r_2, r_3, r_4\}, \{r_{14}, r_{15}, r_{16}\}, -245\}_{A \rightarrow B}$	+30	-35
B^2	$\{\{r_{14}, r_{15}, r_{16}\}, \{r_2, r_3, r_4\}, +133\}_{B \rightarrow A}$	-82	+90
...
A^6	$\{\{r_2, r_3, r_4\}, \{r_{14}, r_{15}, r_{16}\}, -217\}_{A \rightarrow B}$	+2	-7
B^6	$\{\{r_{14}, r_{15}, r_{16}\}, \{r_2, r_3, r_4\}, +223\}_{B \rightarrow A}$	+8	+0
A^1	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).

5. EXPERIMENTAL RESULTS

In order to evaluate and characterize the hypothetical benefit(s) 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 null¹³ 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 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.

¹³For the sake of simplicity, we assumed that a uniform error was made.

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 6.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 3.3.2). Exploring the various other meta-strategies discussed in Section 3.3.2 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). Figure 7.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 6.b presents the mean of the difference in cost of the plans;
- Subjective benefit made by the agents in the deal: Figure 6.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 7.b) and in BR negotiations (Figure 7.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 7.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 al-

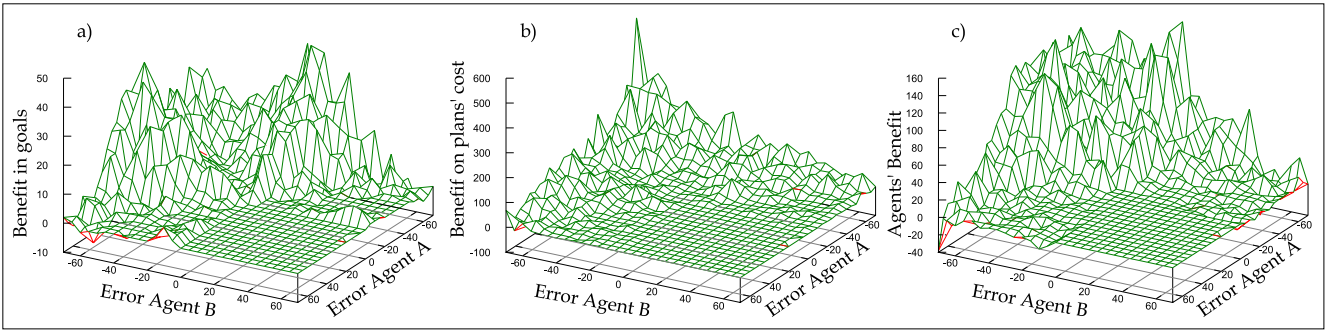


Figure 6: From left to right we have: (a) benefit in term of number of goals achieved and when both achieved their goals we have: (b) benefit on the cost of the plan used, (c) benefit of the agents.




Figure 7: (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.

ternative 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.

5.3 Discussion

The shapes of Figures 6, 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. As a matter of fact, there are a few cases in which Figure 6 indicates 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. The preference ordering for their own plans is defined as follows: $P_1^A \succ_A^c P_2^A \succ_A^c P_3^A$ for A and $P_1^B \succ_B^c P_2^B \succ_B^c P_3^B$ 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' selected plan	Bargaining
P_1^A	P_1^B	failure
P_2^A	P_2^B	failure
P_3^A	P_3^B	success

In the BR case, the first bargaining of this sequence will still fail (like in the BO case). The agents will try to reframe.

Suppose the cost preference ordering of agent B over the plans for agent A is: $P_3^A \succ_B^c P_1^A \succ_B^c P_2^A$. In this case, agent B will suggest that plan P_3^A is actually cheaper than plan P_1^A . Agent A will update his valuation function and take all the plans into account again while agent B will not consider plan P_1^B anymore. The rest of the bargainings sequence will then be:

A's selected plan	B' selected plan	Bargaining
P_3^A	P_2^B	failure
P_1^A	P_3^B	failure
P_2^A	\emptyset	failure

A successful reframing will generally change the combination of bargainings to 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 probability of reaching an agreement.

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: $P_1^A \succ_A^c P_2^A \succ_A^c P_3^A$ and the preference ordering agent B has over the plans of A are: (1) cost preference ordering: $P_2^A \succ_B^c P_3^A \succ_B^c P_1^A$ and (2) profit preference ordering: $P_3^A \succ_B^b P_2^A \succ_B^b P_1^A$. 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 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 A (subjective costs)		
$P_1^A(400)$	$\succ^c P_2^A(500)$	$\succ^c P_3^A(550)$
$P_3^A(475)$	$\succ^c P_1^A(476)$	$\succ^c P_2^A(500)$
$P_1^A(476)$	$\succ^c P_2^A(487.5)$	$\succ^c P_3^A(488.5)$
$P_3^A(481)$	$\succ^c P_1^A(483)$	$\succ^c P_2^A(487.5)$
...		

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 recommendation. Different update functions will also have different impact on this phenomenon.

These two phenomena are a direct consequences of making the agents' preferences dynamic in the context of uncertain and possibly erroneous valuation 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).

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 goals). 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.

The model of automated negotiation proposed in this paper addresses the four limitations attached to previous contributions and mentioned in Section 1:

1. The agents do not have any a-priori knowledge (not even stochastic) of the other's utility function;
2. The reframing 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.
3. 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.
4. The positive results described earlier hold independently of the error made by the agents on the evaluation of the resources not owned.

This robustness to erroneous information about the object of the negotiation, when added to the expressivity gained by not making too strong assumptions, gives a promising ground for building real-world applications in the context of electronic commerce.

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 updating functions, different ways to combine reframing and bargaining and different types of domain structure. 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 then be extended to deal with positive and negative interactions between goals.

Acknowledgments.

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

7. REFERENCES

- [1] H.S. Bierman and L. Fernandez. *Game Theory with Economic Applications*. Addison-Wesley Publishing Company, 2005.
- [2] R. Buttner. A Classification Structure for Automated Negotiations. In *Proceedings of the 4th International Workshop on Multi-Agent Systems in E-Business (MASEB'06)*, pages 523–530. IEEE Computer Society, 2006.
- [3] 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.
- [4] 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.
- [5] V. Lesser, K. Decker, T. Wagner, N. Carver, A. Garvey, B. Horling, D. Neiman, R. Podorozhny, M. NagendraPrasad, A. Raja, R. Vincent, P. Xuan, and X.Q. Zhang. Evolution of the gpgp/taems domain-independent coordination framework. *Autonomous Agents and Multi-Agent Systems*, 9(1):87–143, July 2004. Kluwer Academic Publishers.
- [6] P. B. Linhart, R. Radner, and M.A. Satterthwaite. *Bargaining with incomplete information*. San Diego Academic Press, 1991.
- [7] 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.
- [8] I. Rahwan, S. Ramchurn, N. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation based negotiation. *Knowledge Engineering Review*, 18(4):343–375, 2003.
- [9] 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.
- [10] P Young. *Negotiation Analysis*. University of Michigan Press, 1991.