Bargaining and Argument-based Negotiation:
Some Preliminary Comparisons

Iyad Rahwan[1], Liz Sonenberg[1], and Peter McBurney[2]

[1] Department of Information Systems, University of Melbourne
Parkville, VIC 3010, Australia
i.rahwan@pgrad.unimelb.edu.au, l.sonenberg@unimelb.edu.au
http://www.dis.unimelb.edu.au/agent/
[2] Department of Computer Science, University of Liverpool
Liverpool L69 7ZF, UK
p.j.mcburney@csc.liv.ac.uk

Abstract. Argumentation-based techniques are being increasingly used to construct frameworks for flexible negotiation among computational agents. Despite the advancements made to date, the relationship between argument-based negotiation and bargaining frameworks has been rather informal. This paper presents a preliminary investigation into understanding this relationship. To this end, we present a set of negotiation concepts through which we analyse both bargaining and argumentation-based methods. We demonstrate that if agents have false beliefs, then they may make decisions during negotiation that lead them to suboptimal deals. We then describe different ways in which argument-based communication can cause changes in an agent's beliefs and, consequently, its preferences over contracts. This enables us to demonstrate how the argumentation-based approach can improve both the likelihood and quality of deals.

1 Introduction

Negotiation is a form of interaction in which a group of agents, with conflicting interests and a desire to cooperate, try to come to a mutually acceptable division/exchange of scarce resources [22]. Resources can be commodities, services, time, etc.; in short, anything that is needed to achieve something. Resources are "scarce" in the sense that not all competing claims over them can be simultaneously satisfied.

Frameworks for automated negotiation have been studied analytically using game-theoretic techniques [19] as well as experimentally [4, 6, 10]. Most such negotiation frameworks are focused on *bargaining*, in which the main form of interaction is the exchange of potential deals, i.e., potential allocations of the resources in question.

Recently, it has been proposed that mechanisms for *argumentation* can be used to facilitate negotiation among computational agents. These mechanisms attempt to overcome some of the limitations of bargaining-based frameworks by allowing agents to exchange additional information, or to "argue" about their beliefs and other internal characteristics, during the negotiation process. This process of argumentation allows an agent to *justify* its negotiation stance; and/or *influence* another agent's negotiation stance [9].

Existing literature on argumentation-based negotiation can be roughly classified into two major strands: (i) attempts to adapt dialectical logics for defeasible argumentation by embedding negotiation concepts within these [1, 15, 20]; and (ii) attempts to extend bargaining-based frameworks by allowing agents to exchange rhetorical arguments, such as promises and threats [11, 18].[3]

Despite the advances made to date, the relationship between argument-based negotiation and bargaining frameworks has been rather informal [9]. This paper presents a preliminary investigation into understanding this relationship. To this end, we present a set of negotiation concepts through which we analyse both bargaining and argument-based methods. We demonstrate that if agents have false beliefs, then they may make decisions during negotiation that lead them to suboptimal deals. We then describe different ways in which argument-based communication can cause changes in an agent's beliefs and, consequently, its preferences over contracts. This enables us to demonstrate how the argumentation-based approach can improve both the likelihood and quality of deals.

The paper advances the state of the art in two ways. First, it provides a step towards a more systematic comparison of argument-based and bargaining-based negotiation frameworks. Second, by making the link between belief change and preference change more explicit, we pave the way for the study of negotiation strategies within argument-based frameworks.

The paper is organised as follows. In the next section, we provide a conceptual framework which enables us to capture key negotiation concepts. We use these concepts in section 3 to show how bargaining works and demonstrate how it can lead to suboptimal outcomes. In section 4, we present an abstraction of a class of argument-based negotiation frameworks. We show different ways in which preferences can change due to changes in beliefs, and draw some comparisons with bargaining. We then conclude in section 5.

2   A Conceptual Framework for Negotiation

In this section, we set up the scene for the rest of the paper by formalising the main concepts involved in negotiation.

2.1   Agents and Plans

We have two autonomous agents A and B sharing the same world, which is in some initial state $s \in \mathcal{S}$, where $\mathcal{S}$ is the set of all possible world states. Each agents might, however, believe it is in a different state from s, which can influence its decisions.

To get from one state s_1 to another s_2, agents execute actions. An action $\alpha \in \mathcal{A}$, where $\mathcal{A}$ is the set of all possible actions, moves the world from one state to another; hence it is a function $\alpha : \mathcal{S} \to \mathcal{S}$. We assume that actions are deterministic, and that the world changes only as a result of agents executing actions.[4]

---

[3] For a comprehensive review, the reader may refer to the forthcoming review article [17].

[4] We concede that this treatment of actions is rather simplistic. We made this choice deliberately in order to simplify the analysis.

Definition 1. (Plan) *A one-agent plan* or simply *plan P to move the world from state s_1 to s_2 is a finite list $[\alpha_1, \ldots, \alpha_n]$ of actions such that $s_2 = \alpha_n(\alpha_{n-1}(\ldots \alpha_1(s_1)\ldots))$*

We denote by $\mathcal{P}$ the set of all possible plans. And we denote by $s_1 \models [P]s_2$ that if the world is in state s_1, then executing plan P moves the world to state s_2.

What we have just defined is the *objective* action operators specification, i.e., how the world actually changes as a result of executing actions. Agents, however, might have possibly incomplete or incorrect beliefs about how the world changes as a result of executing actions. We therefore assume each agent i has its own mapping $\alpha^i : \mathcal{S} \cup \{?\} \to \mathcal{S} \cup \{?\}$ for each action, such that always $\alpha^i(?) = ?$. If $\alpha^i_x(s_1) = ?$, then we say that agent i does not know what state action α_x results in if executed in state s_1. The expression $s_1 \models^i [P]s_2$ means that agent i believes executing plan P in state s_1 results in state s_2. Moreover, the expression $s_1 \models^i [P]?$ means that agent i does not know what state results from executing plan P in state s_1.

Agents can evaluate actions and plans based on their costs.

Definition 2. (Cost of Action) *The cost of action α for agent $i \in \{A, B\}$ is defined using an action cost function $Cost : \{A, B\} \times \mathcal{A} \to \mathbb{R}^+$, which assigns a number to each action.*

Definition 3. (Cost of Plan) *The cost of plan $P \in \mathcal{P}$ to agent i is defined using a plan cost function*

$$Cost : \{A, B\} \times \mathcal{P} \to \mathbb{R}^+ \text{ such that } Cost(i, P) = \sum_{\alpha \in P} Cost(i, \alpha)$$

Unlike the case with action operators, where agents can have incorrect beliefs about the results of actions, we assume each agent has accurate knowledge about how much each action costs him/her. However, an agent may not know how much an action would cost another agent (i.e., we only assume each agent i knows accurately what $Cost(i, \alpha)$ is for each α).

Each agent $i \in \{A, B\}$ has a set of desires $\mathcal{D}^i \subseteq \mathcal{D}$, where $\mathcal{D}$ is the set of all possible desires. These desires are formulae in propositional logic or closed formulae in first-order logic (i.e., with no free variables). We say that a world state s satisfies a desire d if $s \models d$, where $\models$ is an appropriate semantic entailment relation.

Definition 4. (Worth of Desire) *The worth of desire d for agent i is defined using a desire worth function $Worth : \{A, B\} \times \mathcal{D} \to \mathbb{R}^+$, which assigns a number to each desire.*

Definition 5. (Worth of State) *The worth of state $s \in \mathcal{S}$ to agent i is defined using a state worth function*

$$Worth : \{A, B\} \times \mathcal{S} \to \mathbb{R}^+ \text{ such that } Worth(i, s) = \sum_{s \models d} Worth(i, d)$$

As with costs, each agent knows precisely what each desire is worth to him/her. Also, an agent may not know how much a desire is worth to another agent (i.e., we only assume each agent i knows accurately what $Worth(i, s)$ is).

We can now define the *utility* of a plan for an agent given it is in a particular state. We distinguish between the *objective* and *perceived* utility. The objective utility denotes the 'actual' gain achieved by the agent based on the actual resulting state (i.e., according to the objective action operators definition). The perceived utility, on the other hand, is the utility the agent 'thinks' it would achieve from that plan, based on what it believes the resulting state is.

Definition 6. (Utility of Plan) *The* utility of plan P *for agent* i *from state* s_1 *is defined as:*

$$Utility(i, P, s_1) = Worth(i, s_2) - Cost(i, P) \text{ where } s_1 \models [P]s_2$$

Definition 7. (Perceived Utility of Plan) *The* perceived utility of plan P *for agent* i *from state* s_1 *is defined as:*

$$Utility^i(i, P, s_1) = Worth(i, s_2) - Cost(i, P) \text{ where } s_1 \models^i [P]s_2$$

Definition 8. (Best Plan) *The* best plan *for agent* i *from state* s_1 *is a plan* $P = BestP(i, s_1)$ *such that* $Utility(i, P, s_1) \geq Utility(i, P', s_1)$ *for all* $P' \neq P$

Definition 9. (Perceived Best Plan) *The* perceived best plan *for agent* i *from state* s_1 *is a plan* $P = BestP^i(i, s_1)$ *such that* $Utility^i(i, P, s_1) \geq Utility^i(i, P', s_1)$ *for all* $P' \neq P$

2.2 Contracts and Deals

So far, we have outlined how an agent can individually achieve its desires through the execution of plans. An agent might also be able to achieve its desires by contracting certain actions to other agents. Since agents are self-interested, they would only perform actions for one another if they receive something in return (i.e., if they get actions done for them, resulting in achieving their own desires). A specification of the terms of such exchange of services is a *contract*.

Definition 10. (Contract) *A* contract Ω *between agents* A *and* B *is a pair* (P_A, P_B) *of plans, and a schedule, such that* P_i, $i \in \{A, B\}$ *is the part of the contract to be executed by agent* i *according to the schedule.*

A schedule is a total order over the union of actions in the two one-agent plans. As with one-agent plans, we denote by $s_1 \models [\Omega]s_2$ that if the world is in state s_1, then executing the contract Ω moves the world to state s_2. Similarly, the perceived result of the contract by agent i is denoted by $s_1 \models^i [\Omega]s_2$. We denote by $\mathcal{C}$ the set of all possible contracts. We now define the cost of a contract to an agent.

Definition 11. (Cost of Contract) *The* cost of contract $\Omega = (P_A, P_B)$ *for agent* $i \in \{A, B\}$ *is the cost of* i*'s part in that contract; i.e.,* $Cost(i, \Omega) = Cost(i, P_i)$

We define the contract's objective and perceived utilities, denoted $Utility(i, \Omega, s_1)$ and $Utility^i(i, \Omega, s_1)$, and the best contract and best perceived contract, denoted $BestC(i, s_1)$ and $BestC^i(i, s_1)$, analogously to plans above.

We can now define the set of contracts acceptable to an agent.

Definition 12. (Individual Rational Contract) *A contract* $\Omega = (P_A, P_B)$ *is individual rational, or simply* acceptable, *for agent i in state s if and only if* $Utility(i, \Omega, s) \geq Utility(i, BestP(i, s), s)$

A *perceived individual rational contracts* is defined similarly using perceived utilities.

A rational agent[5] should only accept contracts that are individual rational. We denote by $IRC(i)$ the set of individual rational contracts for agent i, and by $IRC^i(i)$ the set of perceived individual rational contracts. On this basis, each agent can classify each possible contract into two sets: *acceptable*, *unacceptable*, and *suspended* contracts. Suspended contracts are contracts for which the agent does not know the result (i.e., for which $s_1 \models^i [\Omega]$?), and is hence unable to assess the utilities. If $IRC(i) = \emptyset$, then it makes no sense for agent i to negotiate; i.e., an agent better do things individually.

If agents do not change their beliefs, then the set $IRC^i(i) \cap IRC^j(j)$ is the set of possible deals: contracts that are individual rational from the points of view of both agents. Possible deals are those contracts that make both agents (as far as they know) better off than they would be working individually. If $IRC^i(i) \cap IRC^j(j) = \emptyset$, then agents will never reach a deal unless they change their preferences. Figure 1 exemplifies two cases. Each oval shows the set of individual rational contracts for an agent. If these sets intersect, then a deal is possible.


Fig. 1. Possible and impossible deals

3   Searching for Deals Through Bargaining

In the previous section, we outlined the main concepts involved in the stage prior to negotiation. The questions that raises itself now is the following: *given two agents, each with a set of individual rational contracts, how can agents decide on a particular deal, if such deal is possible?* One way is to search for a deal by suggesting contracts to one another.

---

[5] I.e., rational in the economic sense, attempting to maximise expected utility.

3.1   Elements of Bargaining

Negotiation can be seen as a process of joint search through the space of all contracts (i.e., through the set $\mathcal{C}$), in an attempt to find a mutually acceptable contract (i.e., one that belongs to $IRC^i(i) \cap IRC^j(j)$). Furthermore, agents may wish to find a contract that also satisfies some kind of 'optimality' criteria. For example, agents may attempt to find a contract that is Pareto optimal, or one that maximises the sum or product of their individual utilities.[6]

One of the most widely studied mechanisms for searching for a deal is *bargaining* [12]. In bargaining, agents exchange *offers* – or *proposals*: contracts that represent potential deals. Of course, it would only make sense for each agent to propose contracts that are acceptable to it.

Definition 13. (Offer) *An* offer *is a tuple $\langle i, \Omega \rangle$, where $i \in \{A, B\}$ and $\Omega \in \mathcal{C}$, and represents an announcement by agent i that $\Omega \in IRC^i(i)$.*

During negotiation, each agent may make multiple offers until agreement is reached. At any particular point in time, the offers made constitute the negotiation *position* of the agent: those contracts the agent has announced it is willing to accept as deals. We denote by $\mathcal{O}$ the set of all possible offers (by all agents).[7]

Definition 14. (Position) *The* position *of agent i, denoted $Position(i)$, is a set of contracts i has offered so far, such that at any time, we have $Position(i) \subseteq IRC^i(i)$*

Note that while the set $IRC^i(i)$ is static during bargaining, the set $Position(i)$ is dynamic, since it expands, within the confines of $IRC(i)$, as the agent makes new offers.

A question that raises itself now is: *how does an agent expand its position?* In other words, *given a set of offers made so far, what should an agent offer next?* The answer to this question is what constitutes the agent's bargaining strategy.

Definition 15. (Bargaining Strategy) *A* bargaining strategy *for agent i, denoted Δ^i is a function that takes the history of all proposals made so far, and returns a proposal to make next. Formally: $\Delta^i : 2^{\mathcal{O}} \to \mathcal{O}$, where $2^{\mathcal{O}}$ is the power set of the set of all possible offers $\mathcal{O}$.*

One of the key factors in influencing an agent's negotiation strategy is its preferences over contracts. It would make sense for an agent to begin by offering contracts most preferable to itself, then progressively 'concede' to less preferred contracts if needed.[8] Preference, however, is not the only factor that guides strategy. For example, an agent might have time constraints, making it wish to reach agreement quickly even if such agreement is not optimal. To reach a deal faster, the agent might make bigger concessions than it would otherwise. This issue becomes particularly relevant if the number of possible contracts is very large.

A variety of bargaining strategies have been studied in the literature. Such strategies might be specified in terms of a preprogrammed, fixed sequence of offers [3] or be

---

[6] For more on outcome evaluation, refer to the book by Rosenschein and Zlotkin [19].

[7] Note that $\mathcal{O}$ is different from $\mathcal{C}$. While the latter denotes the set of all possible contracts, the former denotes the set of all possible agent/contract pairs.

[8] This is commonly known as the monotonic concession bargaining strategy.

dependent on factors observed during negotiation itself, such as the offers made by the counterpart [2, 5, 23], or changes in the availability of resources [4]. A thorough examination of these strategies is outside the scope of this study. We note, however, that strategies are highly dependent on the interaction protocol and on the information agents have. For example, following a risk-dependent strategy under the monotonic concession protocol when agents have complete information can be guaranteed to lead to a Pareto-optimal agreement [8]. Such result could not be guaranteed if agents do not know each other's preferences.

3.2   Limitations of Bargaining

One of the main limitations of bargaining frameworks is that they usually assume agents have complete and accurate information about the current world state and the results of actions, and are consequently capable of providing a complete and accurate ranking of all possible contracts. If these assumptions are not satisfied, serious problems start to arise. In particular, bargaining could not be guaranteed to lead to agreements that truly maximise the participants' utilities.

To clarify the above point, consider the following example. Suppose a customer intending to purchase a car assigns a higher preference to Volvos than Toyotas because of his[9] perceived safety of Volvos. Suppose also that this holds despite the customer's belief that Toyotas have cheaper spare parts, because safety is more important to him. If this information is false –for example if Toyota's actually perform as good as Volvos on safety tests–, then the actual utility received by purchasing a Volvo is not maximal. This example is formalised below.

Example 1. Suppose buyer agent B trying to purchase a car from seller A, such that:

- B believes they are in s_1
- $\mathcal{D}^B = \{safety, cheapParts\}$
- $Worth(B, safety) = 18$, $Worth(B, cheapParts) = 12$
- $s_1 \models^B [do_A(giveVolvo), do_B(pay\$10K)]s_2$ where $s_2 \models safety$
- $s_1 \models^B [do_A(giveToyota), do_B(pay\$10K)]s_2'$ where $s_2' \models cheapParts$
- $Cost(B, pay\$10K) = 10$

Then B will assign the following utilities:

- $Utility^B(B, [do_A(giveVolvo), do_B(pay\$10K)], s_1) = 18 - 10 = 8$
- $Utility^B(B, [do_A(giveToyota), do_B(pay\$10K)], s_1) = 12 - 10 = 2$

Consequently, B will attempt to purchase a Volvo. However, suppose that the truth is that:

- $s_1 \models [do_A(giveToyota), do_B(pay\$10K)]s_2''$ where $s_2'' \models cheapParts \wedge safety$

In this case, the actual utility of the Toyota contract would be:

- $Utility(B, [do_A(giveToyota), do_B(pay\$10K)], s_1) = 12 + 18 - 10 = 20$

---

[9] To avoid ambiguity, we shall refer to the seller using she/her and to the buyer using he/his.

Hence, this lack in B's knowledge can lead to negotiation towards a suboptimal deal.

Another case based on the example above is when B does not know about the safety features of cars of make Honda. In this case, B would assign value '?' to Honda contracts, and would be unable to relate it preferentially to Toyotas and Volvos. If Honda's where indeed cheaper, and offer both safety and good spare part prices, agent B would be missing out, again.

What we have just demonstrated is that if agent preferences remain fixed during negotiation and their beliefs are inaccurate, then they may fail to reach deals that maximise their utility. We can generalise this to the following result.

Proposition 1. *In bargaining between agents i and j, the actual best reachable deal is the best deal acceptable to both agents according to their perceived preferences.*

Proof. Let us denote the actual best deal by $BEST(i,j)$. This deal lies in the set $IRC(i) \cap IRC(j)$. But since agents make their decisions based on their perceived contract utilities, each contract $\Omega \notin IRC^i(i) \cap IRC^j(j)$ is unacceptable for at least one agent, and hence will never be selected as a deal. This means that the actual best reachable deal through bargaining is in the set:

$$IRC^i(i) \cap IRC^j(j) \cap IRC(i) \cap IRC(j)$$

Now, if

$$BEST(i,j) \in ((IRC(i) \cap IRC(j)) \backslash (IRC^i(i) \cap IRC^j(j)))$$

then the agents will never reach $BEST(i,j)$. The same thing may apply for the actual second best deal, and so on, until we reach a deal that is within $IRC^i(i) \cap IRC^j(j)$.

This straightforward result demonstrates clearly that as long as agent preferences are inaccurate, they might miss out on better deals. Note, however, that this does not give us an indication of how good or bad the best perceived deal is.

4   Argument-based Negotiation

In the previous section, we explored how bargaining can be used to search for a deal on the basis of fixed agents preferences over contracts. We showed that there are circumstances in which bargaining fails to achieve a deal, or leads to a suboptimal deal. In this section, we explore argument-based approaches to negotiation and relate it to bargaining.

As mentioned earlier, if $IRC^i(i) \cap IRC^j(j) = \emptyset$, then agents will never reach a deal unless at least one of them changes its perceived individually rational contracts. Figure 2 shows two cases where initially no deal was possible because the agents' individual rational contract sets did not intersect, but a deal is enabled by changes in the set of individual rational contracts. In figure 2(a), a deal is enabled when agent B's perceived IRC set changes such that contract X^* becomes acceptable. In figure 2(b), both agents IRCs change, making deal X^{**} (which initially was not acceptable to either agents) mutually acceptable. Changing $IRC^i(i)$ requires changing agent i's preferences, which in

Fig. 2. Changes in perceived individual rational contracts

fact requires change in i's beliefs. *Argumentation* is a way to enable agents to rationally influence such beliefs through rational dialogues.

The benefit of argumentation is apparent in human negotiations. Humans form their preferences based on information available to them. As a result, they acquire and modify their preferences as a result of interaction with the environment and other consumers [13]. Advertising capitalises on this idea, and can be seen a process of argumentation in which marketers attempt to persuade consumers to change their preferences over products [21]. In negotiation, participants are encouraged to argue with one another and discuss each other's interests. This enables them to jointly discover new possibilities and correct misconceptions, which increases both the likelihood and quality of agreement [7]. Computational agents may realise a similar benefit if they are able to conduct dialogues over interests during negotiation.

4.1   Elements of ABN

Argument-based negotiation (ABN) extends bargaining-based protocols. Therefore, concepts such as offers and positions are also part of ABN. In addition, agents can exchange information in order to influence each others' beliefs. As a result, they influence each others' negotiation positions and set of acceptable contracts. The first step towards understanding how preferences over contracts change is, therefore, to understand the different ways influence on beliefs may take place, and how such influence affects the utility an agent assigns to a contract.

Recall that the utility of contracts and plans are calculated by agent i based on the following definition, which merges the definitions of plan and contract utility.

Definition 16. (Utility of Plan or Contract) *The utility of contract or plan X for agent i from state s_1 is defined as: $Utility^i(i, X, s_1) = Worth(i, s_2) - Cost(i, X)$ where $s_1 \models^i [X]s_2$*

From the definition, it is clear that the utility of a contract or plan (a) increases as the the perceived worth of the resulting state increases, and (b) decreases as the perceived cost of carrying out that contract or plan increases. Since we assume that perceived costs are subjective, and are hence accurate, we concentrate on how changes in perceived worth

of state s_2 affect the utility. According to definition 5, the worth of state s_2 depends on the set of desires from $\mathcal{D}^i$ that are satisfied in s_2.

Based on this understanding, we can now enumerate how changes in beliefs can influence the perceived utility of a contract or plan. We dub these changes **C1**, **C2**, etc.

C1 *Learn that in s_1, X results in a state other than s_2*:

   Description: Agent i learns that $s_1 \models^i [P]s_2'$ where $s_2' \neq s_2$.

   Effect: This may trigger a change in the worth of X's result, which then influences the utility of X, as follows:
   1. If $Worth(i, s_2') = Worth^i(s_2)$, then the utility of X remains the same;
   2. If $Worth(i, s_2') \geq Worth(i, s_2)$, then the utility of X increases;
   3. If $Worth(i, s_2') \leq Worth(i, s_2)$, then the utility of X decreases;

   Example: A traveller who knew it was possible to travel to Sydney by train learns that by doing so, he also gets free accommodation with the booking. As a result, his preference for train travel increases. Hence, this is an example of the second effect described above.

C2 *Learn that it is in a different state*:

   Description: The agent learns that it is not in state s_1 as initially thought, but rather in state s_1', where $s_1' \neq s_1$.

   Effect: Two things might happen:
   1. If the agent believes that in this new state, X has the same result, i.e. that $s_1' \models^i [X]s_2'$, then the perceived utility of X remains the same.
   2. If the agent believes X now results in a different state, i.e. that $s_1' \models^i [X]s_2$ where $s_2' \neq s_2$, then the utility of X changes as in the three cases described in **C1** above.

   Example: A traveller who was planning a conference trip learns that the conference has been cancelled. Now, flying to Sydney will no longer achieve his desire to present a research paper.

C3 *Learn a new plan*:

   Description: Agent i, which did not know what plan X results in, i.e., $s_1 \models^i [X]?$, now learns that $s_1 \models^i [X]s_2$.

   Effect: X moves from being suspended to having a precise utility. If X is a contract, it gets classified as either acceptable or unacceptable.

   Example: A car buyer did not know whether a car of make Honda has airbags. After learning that they do, he can now calculate the utility of this car.

C4 *Unlearn an existing plan*:

   Description: Agent i discovers that some X actually does not achieve the expected resulting state, i.e., that $s_1 \models^i [X]?$.

   Effect: The utility of X becomes undefined, and X becomes suspended.

   Example: A traveller might find out that merely booking a ticket does not achieve the state of being in Sydney.

As a result of a perceived utility change, the relative preferences among various plans and contracts *may* change. Preference change may not take place if the agent's perceived utilities of contracts does not change at all, or if utilities do not change enough to cause a reordering of preferences.

Note that what we described above is the effect of a belief change on the utility of a single contract. In fact, each belief change may trigger changes in the utilities of a large number of contracts, resulting in quite complex changes in the agent's preference relation. This adds a significant complexity to strategic reasoning in ABN.

4.2   Embedded Dialogues as Means for Utility Change

One might ask: *on what basis could the above changes in belief and perceived utilities take place during negotiation?* A rational agent should only change its preferences in light of new information. One way to receive such information is through perception of the environment. Another way is through communication with others. Our focus here is on the latter and in particular on situations where belief change happens during the negotiation dialogue itself. In this context, the idea of *embedding* one dialogue in another is relevant. Walton and Krabbe [22, pp. 66] provide a classification of main dialogue types, namely: *persuasion*, *negotiation*, *inquiry*, *deliberation*, *information seeking*, and *eristic* dialogues. Embedding is one type of dialectical shift –moving from one dialogue to another [22, pp. 100–102]. During negotiation between two participants, the following shifts to embedded dialogues may take place:

- *Information seeking in negotiation*: one participant seeks information from its counterpart in order to find out more (e.g., a customer asks a car seller about the safety record of a particular vehicle make);
- *Persuasion in negotiation*: one participant enters a persuasion dialogue in an attempt to change the counterpart's beliefs (e.g., a car salesperson tries to persuade a customer of the value of airbags for safety);
- *Inquiry in negotiation*: both participants initiate an inquiry dialogue in order to find out whether a particular statement is true, or in order to establish the utility of a particular contract; a precondition to enquiry is that neither agent knows the answer a priori (e.g., a customer and car seller jointly attempt to establish whether a particular car meets the customer's safety criteria);
- *Deliberation in negotiation*: both participants enter a deliberation dialogue in order to establish the best course of individual or joint action (i.e., the best plan or joint plan), potentially changing their initial preferences (e.g., a customer and car seller jointly attempt to find out the best way to achieve the customer's safety and budget requirements);

Note that in order to enable the above types of dialogue shifts during negotiation, a protocol that allows dialogue embedding is needed. One such framework was presented by McBurney and Parsons [14].

4.3   Some ABN Examples

We now list a number of examples, building on example 1, which demonstrate some ways in which preference can change as a result of belief change.

Example 2. Car selling agent A initiates the following persuasion dialogue in order to get the buyer B to choose the Toyota:

A: Don't you know that Toyotas actually perform as good as Volvos on major road safety tests?

B: Oh really? And it costs the same right?

A: True.

B: Well, I would rather purchase the Toyota then!

As a result of argumentation, B now believes that

$$s_1 \models^B [do_A(give\,Toyota), do_B(pay\$10K)]s_2'' \text{ where } s_2'' \models cheapParts \wedge safety$$

As we discussed in example 1, this leads to a more accurate preference. Note that this example involves a belief change of type **C1**, where B changes his expectation about the result of the Toyota contract.

Example 3. Suppose B did not initially know about the safety features of cars of make Honda. In this case, B would have the following belief:

$$s_1 \models^B [do_A(giveHonda), do_B(pay\$10K)]?$$

As a result, B would be unable to relate it preferentially to Toyotas and Volvos. Suppose B then initiates the information seeking dialogue:

B: How about that Honda over there?

A: Actually Hondas satisfy both your criteria. They are safe, and also have cheap parts. In fact, this one is available for $8K.

B: Seems better than both. I'll go for the Honda then!

If we have $Cost(B, pay\$8K) = 8$, then as a result of the above dialogue, B can now give a utility valuation for contract $[do_A(giveHonda), do_B(pay\$8K)]$. This will be $12 + 18 - 8 = 22$, which will rank the Honda higher than both Toyotas and Volvos. Note that this example involves a belief change of type **C3** for the Honda contract.

Example 4. Suppose that the seller would still rather sell the Toyota than the Honda, because she wants to get rid of the old Toyota stock. Consider the following dialogue:

B: From what you said, I like this Honda. It offers the same features as the Toyota, but is cheaper.

A: But did you consider its registration cost?

B: It's the same for all cars, so I think it's irrelevant.

A: Actually, the government recently introduced a new tax cut of $3K for purchasing locally manufactured cars. This is aimed at encouraging national industry.

B: Wow! This would indirectly reduce the cost of Toyotas because they are manufactured in Australia. This does not apply to the imported Hondas.

A: That's correct.

B: Aha! Toyota is definitely the way to go then.

Before the dialogue, B knew that if there was a tax cut for local cars, i.e., if it is in $s_1' \models^i localTaxCut$, then purchasing a Toyota results in an additional worth of 3, i.e., that:

$$s_1' \models^i [do_A(give\,Toyota), do_B(pay\$10K)]cheapParts \wedge safety \wedge get\$3K$$

But because B initially thought that there is no such tax cut, i.e., that it is in $s_1 \models^i$ $\neg localTaxCut$, the resulting state was not thought to contain $get\$3K$. During the dialogue B finds out that it is in s_1' rather than s_1. As a result, the utility of the Toyota contract becomes $12 + 18 + 3 - 10 = 23$, whereas the utility of the Honda remains $12 + 18 - 8 = 22$. Note that this dialogue involves a belief change of type **C2**.

4.4   Position and Negotiation Set Dynamics

The examples presented in the previous subsection demonstrate how preferences can change during negotiation as a result of belief and utility changes. Now, the question is: *how can such preference change influence the likelihood and quality of agreement?*

Proposition 2. *Argumentation can influence a negotiator i's set of individually rational contracts.*

This is because changes in utilities may cause existing contracts to leave the set $IRC^i(i)$, or new contracts to enter this set.

Recall from Proposition 1 that the quality of reachable deals depends on the contents of the sets $IRC^i(i)$ (or more specifically, on their intersection) and how they differ from their actual counterparts $IRC(i)$. Hence, changes to $IRC^i(i)$ caused by argumentation could influence the quality of reachable deals. Moreover, argumentation can enable a deal in an otherwise failed negotiation. This happens when the sets of individual rational contracts did not initially intersect.

Proposition 3. *Argumentation can improve the actual quality of the deal reached.*

Proof. Let A and B be two agents negotiating over two mutually acceptable contracts, Ω and Ω'. And suppose that for each agent $i \in \{A, B\}$, the perceived utilities are such that $Utility^i(i, \Omega, s_1^i) \geq Utility^i(i, \Omega', s_1^i)$ whereas actual utilities are such that $Utility(i, \Omega, s_1^i) \leq Utility(i, \Omega', s_1^i)$. This means that contract Ω Pareto dominates[10] Ω' from the point of view of both agents, whereas based on the actual objective utilities, Ω' Pareto dominates Ω. If the agents were bargaining, they would choose Ω. Through argumentation, the beliefs of participants may change such that the perceived utility of Ω' becomes higher than that of Ω for both agents. In this case, Ω' would be chosen, resulting in an objectively better outcome.

A popular example that demonstrates the above proposition has been presented by Parsons et al [15]. The example concerns two home-improvement agents – one trying to hang a mirror, the other trying to hang a painting. They each have some but not all of the resources needed. Even though a deal was possible, the agents could not reach a deal because one agent knew only one way to achieve his goals. By engaging in argument, that agent was able to learn that he could achieve his goals in a different way, by using a different set of resources. Thus, the information exchanged in the course of the interaction resulted in that agent learning a new way to achieve his goal (i.e., learning some new beliefs), and so changed his preferences across the set of possible contracts.

As much as the above result seems promising, there is a flip side to things.

---

[10] I.e., makes one agent better off without making the other worse off.

Proposition 4. *Agents can be worse off as a result of argumentation.*

Proof. Similar to Proposition 3 above, except that the agents begin correctly *preferring Ω', and end up preferring Ω.*

Argumentation can lead to worse outcomes if the resulting preference ordering is more different from the objective ordering than it initially was. Whether and how this happens would depend on the efficiency of the agents' argumentative abilities, their reasoning capabilities and any time constraints, and whether or not they attempt to deceive each other.

5   Conclusions

In this paper, we initiated an investigation into understanding the relationship between bargaining and argumentation-based negotiation frameworks. We described both types of frameworks using a uniform "vocabulary", and made some intuitions more precise. In particular, we provided a precise account of how argumentation can influence preferences over contracts. We then showed how the ability to exchange such arguments can help overcome some problems with bargaining. In particular, we have demonstrated that:

- Rational agents *may* change their preferences in the light of new information;
- Rational agents should *only* change their preferences in the light of new information;
- Negotiation involving the exchange of arguments provides the capability for agents to change their preferences;
- Such negotiations could increase the likelihood and quality of a deal, compared to bargaining, particularly in situations where agents have incomplete and/or inaccurate beliefs;
- Such negotiations could also lead to worse outcomes compared to bargaining;

We are now extending our framework in order to capture richer types of argument-based influences. For example, we are investigating allowing agents to influence each others' desire set itself. In this case, we must distinguish between perceived and actual state worths. The same could be done to plan costs.

Our study also paves the way for a more systematic study of strategies in argument-based negotiation [16]. Understanding the possible effects of different types of embedded dialogues can help an agent make decisions about how to argue during negotiation. This also enables studying more complex strategies that result in multiple related changes in utility. For example, a car seller may first attempt to persuade a customer of adopting a new desire towards safety, then attempt to convince him that his current preferred contract does not achieve this desire.

6   Acknowledgements

References

1. L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the European Conference on Artificial Intelligence (ECAI 2000)*, pages 338–342, Amsterdam, Netherlands, 2000. IOS Press.
2. R. Axelrod. *The Evolution of Cooperation*. Basic Books, Inc. Publishers, New York NY, USA, 1984.
3. A. Chavez and P. Maes. Kasbah: An agent marketplace for buying and selling goods. In *First International Conference on the Practical Appication of Intelligent Agents and Multi-Agent Technology*, 1996.
4. P. Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*. PhD thesis, University of London, Queen Mary and Westfield College, Department of Electronic Engineering, 2000.
5. P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make trade-offs in automated negotiations. *Artificial Intelligence*, 142(2):205–237, 2002.
6. S. Fatima, M. J. Wooldridge, and N. R. Jennings. Multi-issue negotiation under time constraints. In C. Castelfranchi and L. Johnson, editors, *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, pages 143–150, New York, USA, 2002. ACM Press.
7. R. Fisher, W. Ury, and B. Patton. *Getting to Yes: Negotiating Agreement Without Giving In*. Penguin Books, New York, USA, second edition, 1991.
8. J. C. Harsanyi. Approaches to the bargaining problem before and after the theory of games: a critical discussion of Zeuthen's, Hicks', and Nash's theories. *Econometrica*, 24:144–157, 1956.
9. N. R. Jennings, S. Parsons, P. Noriega, and C. Sierra. On argumentation-based negotiation. In *Proceedings of the International Workshop on Multi-Agent Systems*, pages 1–7, Boston, USA, 1998.
10. S. Kraus. *Strategic Negotiation in Multi-Agent Environments*. MIT Press, Cambridge MA, USA, 2001.
11. S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: A logical model and implementation. *Artificial Intelligence*, 104(1–2):1–69, 1998.
12. K. Larson and T. Sandholm. An alternating offers bargaining model for computationally limited agents. In C. Castelfranchi and L. Johnson, editors, *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2002)*, pages 135–142, 2002.
13. G. L. Lilien, P. Kotler, and S. K. Moorthy. *Marketing Models*. Prentice-Hall Press, USA, 1992.
14. P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.
15. S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

16. I. Rahwan, P. McBurney, and L. Sonenberg. Towards a theory of negotiation strategy (a preliminary report). In S. Parsons and P. Gmytrasiewicz, editors, *Proceedings of the 5th Workshop on Game Theoretic and Decision Theoretic Agents (GTDT-2003)*, pages 73–80, 2003.
17. I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation based negotiation. *Knowledge Engineering Review (to appear)*, 2004.
18. S. D. Ramchurn, N. R. Jennings, and C. Sierra. Persuasive negotiation for autonomous agents: a rhetorical approach. In C. Reed, F. Grasso, and G. Carenini, editors, *Proceedings of the IJCAI Workshop on Computational Models of Natural Argument*, pages 9–17. AAAI Press, 2003.
19. J. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, Cambridge MA, USA, 1994.
20. S. V. Rueda, A. J. García, and G. R. Simari. Argument-based negotiation among BDI agents. *Computer Science & Technology*, 2(7), 2002.
21. C. Slade. Reasons to buy: The logic of advertisements. *Argumentation*, 16:157–178, 2002.
22. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, Albany NY, USA, 1995.
23. F. L. B. Zeuthen. *Problems of Monopoly and Economic Warfare*. Routledge and Sons, London, UK, 1930.