A Computational Model of Logic-Based Negotiation

Dongmo Zhang and Yan Zhang
Intelligent Systems Laboratory
School of Computing and Mathematics
University of Western Sydney, Australia
{dongmo,yan}@scm.uws.edu.au

Abstract

This paper presents a computational model of negotiation based on Nebel's syntax-based belief revision. The model guarantees a unique bargaining solution for each bargaining game without using lotteries. Its game-theoretic properties are discussed against the existence and uniqueness of Nash equilibrium and subgame perfect equilibrium. We also study essential computational properties in relation to our negotiation model. In particular, we show that the deal membership checking is DP-complete and the corresponding agreement inference problem is Π_2^P-hard.

Introduction

The methodology of logical frameworks for automated negotiation has received justified attention in recent years(Sycara 1990; Kraus *et al.* 1998; Parsons *et al.* 1998; Booth 2001; Meyer *et al.* 2004; Zhang *et al.* 2004; Zhang 2005). These studies have helped us to establish a qualitative methodology for reasoning about bargaining and negotiation, which differentiates them from the traditional game-theoretic approaches to bargaining problems(Roth 1979; Muthoo 1999). Two different formalisms may be identified in the literature: *argumentation-based frameworks* and *belief-revision-based frameworks*. However, the computational issues of such logic-based formalisms have been largely ignored.

In this paper we develop a bargaining solution based on syntax-dependent formalism of belief revision and discuss its game-theoretic properties and computational properties. We will construct a negotiation function based on Nebel's prioritized belief base revision operation(Nebel 1992). We prove that every possible deal of a bargaining game in our framework is a Nash equilibrium, therefore it is not unique as most game-theory models of bargaining. However, we will show that under certain conditions there exists a unique subgame perfect equilibrium of simultaneous-offers bargaining procedure. Finally we investigate the complexity issue related to the proposed model.

Representation of Bargaining Problems

In this work, we will restrict us to the bargaining situations within which only two agents are involved. We assume that each agent has a set of negotiation items, referred to as *demand set*, which is describable in a finite propositional language $\mathcal{L}$. The language is that of classical propositional logic with an associated consequence operation Cn in the sense that $Cn(X) = \{\varphi : X \vdash \varphi\}$, where X is a set of sentences. A set K of sentences is *logically closed* or called a *belief set* when $K = Cn(K)$. If X, Y are two sets of sentences, $X + Y$ denotes $Cn(X \cup Y)$.

Suppose that X_1 and X_2 are the demand sets of two agents. To simplify exploration, we will use X_{-i} to represent the other set among X_1 and X_2 if X_i is one of them.

Before we present a solution to bargaining problem, we first address the problem of how to represent a bargaining situation.

Preference

In Nash's bargaining theory, a bargaining problem is defined as a pair (S, d) where S is a set of utility pairs that can be derived from feasible agreements and d is a pair of utilities designated to be the "disagreement point". A function that assigns a single outcome to every such problem is a bargaining solution. Notice that the actual negotiation items and their logical relations cannot be explicitly represented in Nash's model. Such an abstraction can be misleading as Rubinstein pointed "*the use of numbers to specify the bargaining problem has obscured the meaning of the Nash bargaining solution*"(Rubinstein 2000). To demonstrate this point, let us consider a simple bargaining situation.

Example 1 (Marriage Contract) *Bob and Mary are going to marry. They negotiate about how to spend their first few years after their marriage. Bob thinks that they should have a child after they marry (c). Mary insists that they should enjoy their life in the first few years (l). Both of them know the fact that they can't have both, i.e.,* $\neg(c \wedge l)$. *Therefore the negotiation items of two persons can be expressed as* $X_1 = \{c, \neg(c \wedge l)\}$ *and* $X_2 = \{l, \neg(c \wedge l)\}$*, respectively.*

We can model the problem either as a two-person non-cooperative game or a Nash bargaining game. Both models require a numerical representation of bargainers' preferences over their negotiation items. Assume that whenever an agreement is reached the winner receives 1 and the loser receives 0; Otherwise, both receive 0. It is easy to see that there are two Nash equilibria in the first game: (c, c) and

(l, l). The outcome of Nash bargaining game is a mixed deal $\frac{1}{2}c + \frac{1}{2}l$. None of the modeling gives a deterministic solution.

We can easily see that the actual utility values in the example do not matter; they just represent bargainers' preferences over their negotiation items. (Osborne and Rubinstein 1990) observe that a pre-ordering over possible outcomes is enough for specifying bargainers' preferences. Accordingly, they define a bargaining problem as a four-tuple $(X, D, \succeq_1, \succeq_2)$, where X is a set of feasible outcomes, D is the disagreement event, and $\succeq_i$ is a complete transitive reflexive ordering representing bargainers's preference relation. Unfortunately such an enhancement of representation does not offer any better solution than Nash's to the above problem even though it gives an explicit representation of feasible agreements. As we will see in Example 3, the negotiation outcome of the problem is in fact determined by the logical relation between negotiation items. (Zhang 2005) defines a bargain problem also as a four-tuple $((X_1, \rho_1), (X_2, \rho_2))$, where X_i is the logical representation of agent i's demand items and ρ_i is a map from a sentence to a real number reflecting bargainer i's entrenchment measurement over possible negotiation items. Unfortunately this definition is not purely logical because numeration is still involved. In this paper, we combine the above two ideas and define a bargaining problem as follows:

Definition 1 *A bargaining game is a four-tuple*[1] $((X_1, \preceq_1), (X_2, \preceq_2))$*, where* X_i *is a logically consistent set of sentences in* $\mathcal{L}$ *and* $\preceq_i$ $(i = 1, 2)$ *is a complete transitive reflexive ordering over* X_i *which satisfies the logical constraint*[2]*:*

(LC) *If* $\varphi_1, \cdots, \varphi_n \vdash \psi$, $\min\{\varphi_1, \cdots, \varphi_n\} \preceq \psi$.

It is easy to show that an ordering satisfying the above conditions uniquely determines an AGM epistemic entrenchment ordering, and vice versa(Gärdenfors 1988). So we will refer to such an ordering as an *entrenchment ordering*.

Hierarchy of negotiation items

Suppose that X is a set of negotiation items from an agent and $\preceq$ the agent's entrenchment ordering over X. We define recursively a hierarchy, $\{X^k\}_{k=1}^{+\infty}$, of X with respect to the ordering $\preceq$ as follows:

1. $X^1 = \{\varphi \in X : \neg\exists\psi \in X(\varphi \prec \psi)\}$;
   $T^1 = X \backslash X^1$.
2. $X^{k+1} = \{\varphi \in T^k : \neg\exists\psi \in T^k(\varphi \prec \psi)\}$;
   $T^{k+1} = T^k \backslash X^{k+1}$.

where $\varphi \prec \psi$ denotes $\varphi \preceq \psi$ and $\psi \not\preceq \varphi$. The intuition behind the construction is that each time collects all maximal elements and remove them from the current set. It is easy to see that there exists a number n such that $X = \bigcup_{k=1}^{n} X^k$.

In the sequent, we write $X^{\leq l}$ to denote $\bigcup_{k=1}^{l} X^k$.

[1] More precisely a pair of pairs.

[2] The condition of LC is introduced by (Zhang 2005).

Let O be any set of sentences in $\mathcal{L}$, we define the *degree of coverage* of O over X, denoted by $\rho_X(O)$, to be the greatest number l which satisfies $X^{\leq l} \subseteq O$.

Prioritized base revision

Once we have a hierarchy of negotiation items of each agent, we will be able to define a belief revision function for each agent, which will play a central role in our construction of bargaining solution. Following (Nebel 1992), we define a *prioritized base revision function* $\otimes$ as follows:

For any sets X and F of sentences and an entrenchment ordering $\preceq$ over X,

$$X \otimes F \stackrel{def}{=} \bigcap_{H \in X \Downarrow F} Cn(H) + F.$$

where $X \Downarrow F$ is defined in the following: $H \in X \Downarrow F$ if and only if

1. $H \subseteq X$,
2. for all $k = 1, 2, \cdots$, $H \cap X^k$ is set-inclusion maximal among the subsets of X^k such that $\bigcup_{j=1}^{k} (H \cap X^j) \cup F$ is consistent.

Lemma 1 [Nebel 1992] *If* X *is logically closed and* $\preceq$ *is an entrenchment ordering over* X*, then* $\otimes$ *satisfies all AGM postulates.*

Bargaining Solution

According to Nash's definition, a bargaining solution is a function that assigns a single outcome to every bargaining game. In this section we will construct such a function in relation to our definition of bargaining games. First, let us consider all the possible outcomes of a negotiation game.

Possible deals

A deal is the concessions made by two negotiating agents. Such a deal can be defined as a pair of subsets of two agents' demand sets. Considering the real-life bargaining, bargainers normally intend to keep their highly entrenched negotiable items and give up those less entrenched items if necessary. This idea leads to our definition of deals based on agents' hierarchy of negotiation items.

Definition 2 *Let* $B = ((X_1, \preceq_1), (X_2, \preceq_2))$ *be a bargaining game. A deal of* B *is a pair* (D_1, D_2) *satisfying the following two conditions: for each* $i = 1, 2$,

1. $D_i \subseteq X_i$;
2. *for each* $k = 1, 2, \cdots$, $D_i \cap X_i^k$ *is set-inclusion maximal among the subsets of* X_i^k *such that* $\bigcup_{j=1}^{k} (D_i \cap X_i^j) \cup D_{-i}$ *is consistent.*

The set of all deals of B *is denoted by* $\Omega(B)$.

It is easy to see that if $X_1 \cup X_2$ is consistent, $\Omega(B) = \{(X_1, X_2)\}$.

Example 2 *Consider all the possibilities of each person's preferences with the scenario described in Example 1. The possible deals for each case will then be the following:*

Games	*Bob's pref.*	*Mary's pref.*	*Deals*
B^1	$\neg(c \wedge l) \succ_1 c$	$l \succ_2 \neg(c \wedge l)$	D^1, D^2
B^2	$c \succ_1 \neg(c \wedge l)$	$\neg(c \wedge l) \succ_2 l$	D^1, D^2
B^3	$\neg(c \wedge l) \succ_1 c$	$\neg(c \wedge l) \succ_2 l$	D^1, D^2
B^4	$c \succ_1 \neg(c \wedge l)$	$l \succ_2 \neg(c \wedge l)$	D^1, D^2, D^3

where
$D^1 = (\{\neg(c \wedge l), c\}, \{\neg(c \wedge l)\})$;
$D^2 = (\{\neg(c \wedge l)\}, \{l, \neg(c \wedge l)\})$;
$D^3 = (\{c\}, \{l\})$.

The core of agreement

Example 2 shows that in each game, the set of possible deals lists all possible agreements the negotiation could reach. However, it is left uncertain that which deal will be most likely to be the final agreement. Before we answer the question, let us analyze which items are most likely to be included in the final agreement if the negotiation procedure is "*fair*" to each individual.

Given a bargaining game B and a deal $D = (D_1, D_2)$ of the game, let

$$\rho_B(D) \stackrel{def}{=} \min\{\rho_{X_1}(D_1), \rho_{X_2}(D_2)\}$$

$$\rho_B \stackrel{def}{=} \max\{\rho_B(D) : D \in \Omega(B)\}$$

We call $\rho_B(D)$ the *degree of coverage of deal D* and ρ_B the *degree of coverage of the game*.

Let $\gamma(B) = \{D \in \Omega(B) : \rho_B(D) = \rho_B\}$, representing the subset of $\Omega(B)$ that contains the deals with the highest degree of coverage over all deals in $\Omega(B)$. And let

$$\Phi_1 \stackrel{def}{=} \bigcap_{(D_1,D_2)\in\gamma(B)} D_1, \quad \Phi_2 \stackrel{def}{=} \bigcap_{(D_1,D_2)\in\gamma(B)} D_2$$

We call $\Phi = (\Phi_1, \Phi_2)$ the *core of the game*. The min-max construction of the core captures the idea that the final agreement should maximally and evenly satisfy both agents's demands.

It is easy to verify that $\rho_B = \min\{\rho_{X_1}(\Phi_1), \rho_{X_2}(\Phi_2)\}$.

Bargaining solution

Now we can finalize the construction of our bargaining model.

Definition 3 *A bargaining solution is a function* $\mathbf{A}$ *which maps a bargaining game to a set of sentences (agreement), defined as follows. For each bargaining game* $B = ((X_1, \preceq_1), (X_2, \preceq_2))$

$$\mathbf{A}(B) \stackrel{def}{=} (X_1 \otimes_1 \Phi_2) \cap (X_2 \otimes_2 \Phi_1) \quad (1)$$

where (Φ_1, Φ_2) *is the core of* B *and* $\otimes_i$ *is the prioritized base revision operator over* $(X_i, \preceq_i)$.

There are a few points we would like to make here:

1. A bargaining process takes two stages. In the first stage, each agent agrees on accepting the other party's core demands if the other agent does the same. In the second stage, each agent adjust its own demands in order to make them consistent with the reached agreement in the first stage by conducting a course of belief revision.
2. Different from most existing belief-revision-based frameworks for negotiation, the bargaining model we define above is *syntax dependent*. However, if we restrict the bargaining games to be the ones where the demand sets are logically closed, the bargaining solutions will be syntax independent. A detailed discussion on the relationship between the existing approaches and the current work will be given in the full version of the paper.
3. Unlike Nash's definition, our bargaining solution assigns a set of sentences to each bargaining game. In other words, the bargaining solution directly gives the actual contract resulted from the bargaining rather than a pair of concessions made by two agents.
4. Similar to (Zhang *et al.* 2004), the outcome of bargaining is defined as the intersection of two agents' revision sets. However, the negotiation function they define is based on generic AGM belief revision operators without concrete construction. No computational model is given.
5. Different from Nash's solution and (Zhang 2005)'s solution, we do not require to play a lottery (mixed deals) to settle a tie bargaining situation. The following example will demonstrate how this can happen.

Example 3 *Consider the Marriage Contract example again. The solution of each bargaining game defined in Example 2 is respectively:*
$A(B^1) = Cn(\{\neg c, l\})$; $A(B^2) = Cn(\{c, \neg l\})$;
$A(B^3) = Cn(\{\neg(c \wedge l)\})$; $A(B^4) = Cn(\{c, l\})$.

In the first game B^1, Bob is more level-headed who entrenches the commonsense ($\neg(c \wedge l)$) more than his personal demand (c), which causes him losses the game. B^2 is just the opposite to B^1. In the third game, nobody wins the game. However, since both of them are quite rational they might settle down to face the reality. The last game reaches a "both-win" situation because none of them cares too much about what the reality is (Hopefully they can create another reality).

The example shows that logical reasoning plays a key role in negotiation and bargaining. The logic-based formalism of bargaining can offer a more subtle solution than game-theoretic approaches. Our model gives a deterministic solution to each tie situation (game 3 and 4). Note that our solution does not necessarily satisfy *Pareto Optimality* because in the tie situations we require both players make concession rather than randomly pick up one to do that by using a lottery. In fact, this example provides an intuitive counter-example against this controversial requirement of Nash's solution (Roth 1979).

The following theorem guarantees the final agreement contains the consents reached in the first stage of a bargaining procedure.

Theorem 1 *For any bargaining game* B, $\Phi_1 + \Phi_2 \subseteq \mathbf{A}(B)$.

Proof: We only prove $\Phi_1 \subseteq X_1 \otimes_1 \Phi_2$. The rest of the proof is straightforward. According to the definition of prioritized base revision, we have $X_1 \otimes_1 \Phi_2 = \bigcap_{H \in X_1 \Downarrow \Phi_2} Cn(H) + \Phi_2$. For any $H \in X_1 \Downarrow \Phi_2$, there is a deal $(D_1, D_2) \in \Omega(B)$ such that $D_1 = H$. This is because we can extend the pair (H, Φ_2) to a deal (H, D_2) so that D_2 satisfies the conditions in Definition 2. On the other hand, since $X^{\leq \rho_B} \cup \Phi_2$ is consistent, we have $\rho_{X_1}(H) \geq \rho_B$. Therefore $(D_1, D_2) \in \gamma(B)$. Since $\Phi_1 \subseteq D_1$, we then have $\Phi_1 \subseteq H$. We conclude that $\Phi_1 \subseteq X_1 \otimes_1 \Phi_2$. ¶

Game-Theoretic Properties

The model we presented in the previous section has most desired logical properties and game-theoretic properties. Due to space limit, we omit the presentation of logical properties and concentrate on its game-theoretic properties. We assume that readers are familiar with the basic concepts in game theory, such as extensive form of a game, Nash equilibrium, and subgame perfect equilibrium (Osborne and Rubinstein 1990).

Utilities and Nash equilibrium

In equilibrium analysis, two concepts play an essential role: *strategy* and *utility*. Given a bargaining game $B = ((X_1, \preceq_1), (X_2, \preceq_2))$, a *strategy profile* of the game is a pair (S_1, S_2) where $S_1 \subseteq X_1$ and $S_2 \subseteq X_2$. The intended meaning for a strategy is that during bargaining process each player places a subset of her demand set as her proposal toward an agreement.

A strategy profile is called to be *compatible* if it satisfies:

1. $S_1 \cup S_2$ is consistent;
2. $S_1 \subseteq X_1 \otimes_1 S_2$ and $S_2 \subseteq X_2 \otimes_2 S_1$.

We assume that a strategy profile can lead to an agreement if and only if it is compatible. The outcome of the game will be $(X_1 \otimes_1 S_2) \cap (X_2 \otimes_2 S_1)$. Notice that in this case $S_1 \cup S_2 \subseteq (X_1 \otimes_1 S_2) \cap (X_2 \otimes_2 S_1)$, each player's demand has been met.

Now we can define the utility for each game player. If a game ends with an agreement, we define player $i's$ payoff is the degree of coverage of her strategy, i.e., $\rho_{X_i}(S_i)$. If the game ends with disagreement, then each player's payoff is zero.

Theorem 2 *For any bargaining game B, each deal of the game is a Nash equilibrium.*

Proof: Let $(D_1, D_2) \in \Omega(B)$. It is easy to verify that (D_1, D_2) is a compatible strategy profile. Now we show that if one player, say player 2, use the strategy D_2, then D_1 is optimal to player 1 w.r.t her payoff. Suppose otherwise there is a strategy S_1 for player 1 such that $\rho_{X_1}(S_1) > \rho_{X_1}(D_1)$. We then have $X_1^{\leq \rho_{X_1}(S_1)} \subseteq S_1$. Thus $\bigcup_{j=1}^{\rho_{X_1}(S_1)} (D_1 \cap X_1^j) \cup D_2$ is not set-inclusion maximal, which contradicts the definition of deals. Therefore player 1 has no incentive to vary her strategy. ¶

The result shows that the concept of Nash equilibrium is too weak to determine a bargaining solution, which is similar to the result in game theory(see (Osborne and Rubinstein 1990) p. 41). Interestingly, this theorem provides a natural link between the essential AI approach - *minimization* - and the fundamental game theory approach - *equilibrium analysis*.

Simultaneous-offers model of bargaining

In order to get a more refined concept of equilibrium, we devise a bargaining procedure with multiple simultaneous offers from each player.

Assume that the bargaining between two agents follows the following simultaneous-offers procedure. Given a bargaining game $B = ((X_1, \preceq_1), (X_2, \preceq_2))$, at each period, two players simultaneously make a proposal of demands: (S_1, S_2), where $S_1 \subseteq X_1$ and $S_2 \subseteq X_2$. If (S_1, S_2) is compatible, then an agreement is reached. Otherwise, the game continues for at least one more period, and each player has an opportunity to make a new proposal by reducing at least one item from her previous demand. In other words, each player should delete at least one sentence (as well as its logical equivalent sentences) from her previous proposal. A player could also choose to stand still by holding her previous demand. The game ends with agreement at any period in which the demands of the two players are compatible; or else the game ends with disagreement at any period in which both players choose to stand still prior to reaching agreement, in this case no agreement is reached.

Subgame perfect equilibrium

Given a bargaining game $B = ((X_1, \preceq_1), (X_2, \preceq_2))$, let $\Phi_1^* = X_1 \cap (X_1 \otimes_1 \Phi_2)$ and $\Phi_2^* = X_2 \cap (X_2 \otimes_2 \Phi_1)$, where (Φ_1, Φ_2) is the core of B. We call (Φ_1^*, Φ_2^*) the *completion* of (Φ_1, Φ_2).

Theorem 3 *If both X_1 and X_2 are logically closed and $\Phi_1^* \cup \Phi_2^*$ is consistent, then (Φ_1^*, Φ_2^*) is a subgame perfect equilibrium(SPE) of the simultaneous-offers model*[3]. *Moreover, (Φ_1^*, Φ_2^*) is the unique SPE in the sense that for any compatible strategy profile (S_1, S_2) such that $\Phi_1 \subseteq S_1$ and $\Phi_2 \subseteq S_2$, $\rho_{X_i}(S_i) \leq \rho_{X_i}(\Phi_i^*)$ $(i = 1, 2)$.*

Proof: To prove the theorem, we need the following technical lemma: *If X_1 and X_2 are belief sets and $\Phi_1^* \cup \Phi_2^*$ is consistent, then for each $i = 1, 2$,*

1. $\Phi_i^* \subseteq X_i \otimes_i \Phi_{-i}^*$;
2. $X_i \otimes_i \Phi_{-i} = X_i \otimes_i \Phi_{-i}^*$;
3. $\rho_i(\Phi_i^*) = \rho_i(X_i \otimes_i \Phi_{-i}^*)$.

Now we prove that (Φ_1^*, Φ_2^*) is a SPE. First, at any period of the game, if the strategy (Φ_1^*, Φ_2^*) is used, an agreement will be reached and the payoff of the players will be $(\rho_{X_1}(\Phi_1^*), \rho_{X_2}(\Phi_2^*))$. Assume that only one player, say player 1, applies the strategy and player 2 instead places a proposal S_2. Whenever an agreement is reached, (Φ_1^*, S_2) must be compatible, i.e.,

1. $\Phi_1^* \cup S_2$ is consistent,
2. $\Phi_1^* \subseteq X_1 \otimes_1 S_2$ and $S_2 \subseteq X_2 \otimes_2 \Phi_1^*$.

The agreement will then be $(X_1 \otimes_1 S_2) \cap (X_2 \otimes_2 \Phi_1^*)$ and the payoffs will be $(\rho_{X_1}(\Phi_1^*), \rho_{X_2}(S_2))$. By above condition 2, $\rho_{X_2}(S_2) \leq \rho_{X_2}(X_2 \otimes_2 \Phi_1^*) = \rho_{X_2}(\Phi_2^*)$. Therefore deviation to S_2 from Φ_2^* is not profitable for player 2. If no agreement is reached, which means (Φ_1^*, S_2) is incompatible, the game will

[3] A strategy profile is a subgame perfect equilibrium of an extensive form of a game if the strategy pair it induces in every subgame is a Nash equilibrium of that subgame (see (Osborne and Rubinstein 1990) p43).

move to next round. As assumed, player 1 will play Φ_1^* again. If player 2 also stand still, then the game ends with disagreement in which player 2 does not achieve any better. On the other hand, if player 2 chooses to reduce her demand, the game continues with a situation similar to above but could be even worse (never be better because condition 2 needs to be true whenever an agreement is reached).

Finally we prove the uniqueness. Since (S_1, S_2) is compatible, we have $S_i \subseteq X_i \otimes_i S_{-i}$. Thus $\rho_{X_i}(S_i) \leq \rho_{X_i}(X_i \otimes_i S_{-i}) \leq \rho_{X_i}(X_i \otimes_i \Phi_{-i}) = \rho_{X_i}(\Phi_i^*)$. ¶

Note that both conditions for the existence and uniqueness of SPE are crucial. On one hand, we requires that the demand sets are logically closed because the belief revision operator we use to construct the negotiation function is syntax-dependent. On the other hand, if the completion pair of the core of a game is inconsistent, there could be multiple SPEs. For instance, suppose $X_1 = \{p\}$ and $X_2 = \{\neg p\}$. Then the core completion is inconsistent. In such a case, we can easily to verify that there are two SPEs, both of which contain the core. Nevertheless, the result of Theorem 3 is still significant because the simultaneous-offers model of bargaining in game-theoretic bargaining theory could have even continuum SPE (see (Muthoo 1999) p191).

Computational Properties

In this section, we study the computational properties of the negotiation model that we developed earlier. We assume that readers are familiar with the complexity classes of P, NP, coNP, Σ_2^P and Π_2^P =coΣ_2^P. The class of DP contains all languages L such that $L = L_1 \cap L_2$ where L_1 is in NP and L_2 is in coNP. It is well known that $P \subseteq NP \subseteq DP \subseteq \Sigma_2^P$, and these inclusions are generally believed to be proper (readers may refer to (Papadimitriou 1994) for further details).

Consider a bargaining game $B = ((X_1, \preceq_1), (X_2, \preceq_2))$ where X_1 and X_2 are finite. From the hierarchy definition on X_1 and X_2, it is clear that for each $i = 1, 2$, we can always write $X_i = X_i^1 \cup \cdots \cup X_i^m$, where $X_i^k \cap X_i^l = \emptyset$ for any $k \neq l$. Also for each $k < m$, if a formula $\varphi \in X_i^k$, then there does not exist a $\psi \in X_i^l$ $(k < l)$ such that $\varphi \prec_i \psi$. Therefore, for the convenience of our complexity analysis, in the rest of this section, we will specify a bargaining game as $B = (X_1, X_2)$, where $X_1 = \bigcup_{i=1}^m X_1^i$ and $X_2 = \bigcup_{j=1}^n X_2^j$, and $X_1^1, \cdots, X_1^m$, and $X_2^1, \cdots, X_2^n$ are the partitions of X_1 and X_2 respectively and satisfy the property mentioned above. Sometimes we may simply write $B = (X_1, X_2)$ wihout specifying partitions on X_1 and X_2 if they are not important in our discussion. Under our notation, the definition of agreement function can be simplified as $\mathbf{A}(B) = (X_1 \otimes \Phi_2) \cap (X_2 \otimes \Phi_1)$, where we do not distinguish operators $\otimes_1$ and $\otimes_2$ (in the original definition) which correspond to the entrenchment orderings $\preceq_1$ and $\preceq_2$, respectively.

Theorem 4 *Let $B = (X_1, X_2)$ be a bargaining game, and $D_1 \subseteq X_1$ and $D_2 \subseteq X_2$. Deciding whether (D_1, D_2) is a deal of B is DP-complete.*

Proof: Membership proof. According to Definition 2, to decide whether (D_1, D_2) is a deal of B, for D_1 (or D_2), we need to check: (1) for each $k = 1, \cdots, m$ (or for $k' = 1, \cdots, n$ resp.), whether $D_2 \cup \bigcup_{j=1}^{k} (D_1 \cap X_1^j)$ (or $D_1 \cup \bigcup_{j=1}^{k'} (D_2 \cap X_2^j)$ resp.) is consistent; and (2) such D_1 and D_2 are maximal such subsets of X_1 and X_2 respectively. Clearly, for each k, the set $\bigcup_{j=1}^k (D_1 \cap X_1^j)$ can be computed in polynomial time, and checking the consistency of $D_2 \cup \bigcup_{j=1}^{k} (D_1 \cap X_1^j)$ is in NP. The same for D_2 case. In order to check whether D_1 and D_2 are the maximal subsets of X_1 and X_2 respectively satisfying the condition, for each $\varphi \in (X_1 \setminus D_1)$ and $\psi \in (S_2 \setminus D_2)$, for each k, we check the inconsistency of $D_2 \cup \{\psi\} \cup \bigcup_{j=1}^{k} ((D_1 \cup \{\varphi\}) \cap X_1^k)$ (the same for D_2 as well). There are $(|X_1| - |D_1|) \cdot (|X_2| - |D_2|)$ such $(D_1 \cup \{\varphi\}, D_2 \cup \{\psi\})$ to check, which can be done in coNP. So the problem is in DP.

Hardness proof. It is known that for given propositional formulas φ_1 and φ_2, deciding whether φ_1 is satisfiable and φ_2 is unsatisfiable is DP-complete (Papadimitriou 1994). Given two propositional formulas φ_1 and φ_2, we construct in polynomial time a transformation from the φ_1's satisfiability and φ_2's unsatisfiability to a deal decision problem of a game. We simply define a game $B = (X_1, X_2) = (\{a, a \supset \varphi_2\}, \{\varphi_1\})$, where a is a new propositional atom not occurring in φ_1 or φ_2, and no partition is specified on X_1 or X_2. Let $D_1 = X_1 \setminus \{a \supset \varphi_2\}$ and $D_2 = X_2 = \{\varphi_1\}$. Clearly, D_1 and D_2 are the maximal subsets of X_1 and X_2 respectively such that $D_1 \cup D_2$ is consistent iff φ_1 is satisfiable and φ_2 is unsatisfiable. ¶

Now we study computational issues in relation to the bargaining solution function (Definition 3). As described earlier, the degree of coverage of deals plays an essential role in the definition.

Proposition 1 *Given a bargaining game $B = (X_1, X_2)$ and a deal (D_1, D_2) of the game. The degree of coverage of (D_1, D_2) can be computed in time $\mathcal{O}(|X_1| + |X_2|)$.*

From Proposition 1 we know that computing a deal's degree of coverage is easy. However, the following result implies that finding the core of agreement for a given game is quite difficult.

Theorem 5 *Given a bargaining game B and a pair of sets of formulas $\Phi = (\Phi_1, \Phi_2)$. Deciding whether Φ is the core of B is NP-hard as well as coNP-hard.*

Proof: We consider a special case that $\gamma(B)$ only contains one deal of B. Then our problem reduces to decide whether the given pair of sets of formulas (now we write them as (D_1, D_2) is in $\gamma(B)$). We show this decision problem is NP-hard as well as coNP-hard. Due to a space limit, we only present the proof of coNP hardness, while NP hardness is proved in a similar style but somewhat more tedious.

Given a propositional formula φ whose set of variables is $V = \{x_1, \cdots, x_n\}$. Based on φ, we construct in polynomial time a bargaining game $B = (X_1, X_2)$ and specify a pair (D_1, D_2) where $D_1 \subseteq X_1$ and $D_2 \subseteq X_2$, and show that φ is valid iff (D_1, D_2) is a deal of B with a maximal degree of coverage. Let $\{a, b, c, p, p_1, p_2, p_3, p_4\}$ be newly introduced variables not occurrings in φ. We specify a game $B = (X_1 = \bigcup_{i=1}^{5} X_1^i, X_2 = \bigcup_{j=1}^{5} X_1^j)$, where $X_1 = X_1^1 \cup X_1^2 \cup X_1^3 \cup X_1^4 \cup X_1^5 = \{p\} \cup$

$\{\neg a\} \cup \{\neg b\} \cup \{\neg x_1, \cdots, , \neg x_n\} \cup \{\neg c\}$, and $X_2 = X_2^1 \cup X_2^2 \cup X_2^3 \cup X_2^4 \cup X_2^5 = \{\mu\} \cup \{p_1\} \cup \{p_2\} \cup \{p_3\} \cup \{p_4\}$, and $\mu = [(\varphi \supset (a \wedge b)) \wedge (\bigwedge \neg x_i \wedge c)] \vee [\neg\varphi \wedge \neg a \wedge b \wedge \neg c]$. Clearly, both X_1 and X_2 are consistent but $X_1 \cup X_2$ is not consistent. Now we specify $D_1 = \{p, \neg x_1, \cdots, \neg x_n\}$ and $D_2 = X_2 = \{\mu, p_1, p_2, p_3, p_4\}$. We will show that φ is valid iff (D_1, D_2) is a deal of B with a maximal degree of coverage 1. That is, $(D_1, D_2) \in \gamma(B)$.

Clearly, if φ is valid, $[\neg\varphi \wedge \neg a \wedge b \wedge \neg c]$ cannot be satisfied. So the only deal of B is (D_1, D_2) which is of the degree of coverage 1 (note that $X_1^1 \subseteq D_1$ and $X_1^1 \cup X_1^2 \not\subseteq D_1$). On the other hand, if φ is not valid, it means that formula $[\neg\varphi \wedge \neg a \wedge b \wedge \neg c]$ is satisfiable. In this case, there exists a deal (D_1', D_2) of B where $\{p, \neg a, \neg b\} \subseteq D_1'$ such that $D_1' \cup D_2$ consistent. This implies that the degree of coverage of (D_1', D_2) equals to or is greater than 3 as $X_1^1 \cup X_1^2 \cup X_1^3 \subseteq D_1'$. So $(D_1, D_2) \notin \gamma(B)$. ¶

Given a bargaining game B and an agreement function $\mathbf{A}$, the inference problem of negotiation is to decide whether some φ is derivable from $\mathbf{A}(B)$. This is also difficult, even if we restrict B to a very special case.

Theorem 6 *Let B be a bargaining game and φ a formula. Deciding whether $\mathbf{A}(B) \vdash \varphi$ is Π_2^P-hard.*

Proof: We consider a special $B = (X_1, X_2)$, where no partitions are put on X_1 and X_2, and $X_2 = \{\top\}$. It is easy to see that each deal in $\gamma(B)$ is of the form $(D_1, \{\top\})$ where D_1 is a maximal consistent subset of X_1. In this case, $\Phi_1 = \bigcap_{(D_1, \{\top\}) \in \gamma(B)} D_1$ and $\Phi_2 = \{\top\}$. Therefore, $\mathbf{A}(B) = (X_1 \otimes \{\top\}) \cap (\{\top\} \otimes \Phi_1) = X_1 \otimes \{\top\}$. So we have $\mathbf{A}(B) \vdash \varphi$ iff $X_1 \otimes \{\top\} \vdash \varphi$. Clearly, deciding $X_1 \otimes \{\top\} \vdash \varphi$ is Π_2^P-complete (Nebel 1992). So deciding $A(B) \vdash \varphi$ is Π_2^P-hard. ¶

From Theorem 6, we can see that the syntax based negotiation is at least as hard as Nebel's syntax based belief revision. From Theorem 5, we further observe that the core of agreement checking for B is probably beyond $NP \cup coNP$ (otherwise it will result in NP=coNP). Since computing the core of agreement is a precondition to compute $\mathbf{A}(B)$, we believe that the inference problem for syntax based negotiation is probably beyond Π_2. It follows that syntax based negotiation is generally harder than syntax based belief revision. Although this result is not encouraged from a computational viewpoint, it indeed informs that the logic based negotiation is a more complex process than belief revision.

The following proposition states that under the condition that $\gamma(B)$ is provided, the inference problem for negotiation is in Π_2^P.

Proposition 2 *Let B be a bargaining game and φ a formula. Deciding whether $\mathbf{A}(B) \vdash \varphi$ is in Π_2^P provided that $\gamma(B)$ is given.*

Conclusion and Related work

In this paper, we have developed a logical framework of bargaining based on Nebel's prioritized belief base revision. In the framework, a negotiation procedure is divided into two stages. In the first stage, two agents meet together to work out mutually acceptable core demand from each side. In the second stage, each agent adjusts its own demand in order to make it consistent with the consent reached in the first stage by conducting a course of belief revision. The outcome of the bargaining is then the intersection of two agents' revised demand sets. We have proved that in certain conditions the pair of revised demand sets determines a unique subgame perfect equilibrium of simultaneous-offers model. Our complexity analysis indicates that the computation related to a syntax-based negotiation is more difficult than syntax-based belief revision.

There are several pieces of work in the literature which use belief revision as a tool to build negotiation model. (Booth 2001; Meyer *et al.* 2004; Zhang *et al.* 2004) presented a series of AGM-like frameworks of negotiation operation. These frameworks are all syntax-independent and deal with only generic properties of negotiation processes. Additionally, since there is no representation of players' preferences in these frameworks, it is unclear how these modelings are related to the game-theoretic approaches. (Zhang 2005) proposed a bargaining solution which combines Nash's axiomatic approach and belief-revision-based approach with a real-number representation of player's preferences. Since their definition of bargaining solution relies on the representation and calculation of real numbers, their solution cannot be expressed purely in logic. More significantly, the uniqueness of their bargain solution is not guaranteed without introducing mixed deals. Finally, none of the above approaches addressed the computational issue of logic-based negotiation in detail[4].

References

R. Booth, 2001. A negotiation-style framework for non-priorised revision, *TARK'01*, 137-150.

P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States,* The MIT Press.

S.Kraus, K.Sycara, A.Evenchik, 1998. Reaching agreements through argumentation: a logical model and implementation, *Artificial Intelligence*, 104(1998), 1-69.

T. Meyer, N. Foo, R. Kwok and D. Zhang, 2004b. Logical foundations of negotiation: strategies and preferences, *KR'04*, 311-318.

A. Muthoo, 1999. *Bargaining Theory with Applications*, Cambridge University Press.

B. Nebel, 1992. Syntax-based approaches to belief revision, in: P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, 52-88.

M. J. Osborne and A. Rubinstein, 1990. *Bargaining and Markets*, Academic Press.

S. Parsons, C. Sierra, N. Jennings, 1998. Agents that reason and negotiate by arguing, J. of Logic and Computation 8(3): 261-292.

C.H. Papadimitriou, 1994. *Computational Complexity*. Addison Wesley.

A. Roth, 1979. *Axiomatic Models of Bargaining*, Springer-Verlag.

K. Sycara, 1990. Persuasive argumentation in negotiation, *Theory and Decision*, 28:203-242.

A. Rubinstein, 2000. *Economics and Language: Five Essays*, Cambridge University Press.

D. Zhang, N. Foo, T. Meyer and R. Kwok, 2004. Negotiation as mutual belief revision, *AAAI-04*, 317-322.

D. Zhang, 2005. A logical model for Nash bargaining solution, *IJCAI-05* ,983-988.

[4]Belief-revision-based negotiation has a close relation with belief merging. Due to space limit, we leave the discuss to the full version of the paper.