

# Logical foundations of negotiation: Outcome, Concession and Adaptation

**Thomas Meyer and Norman Foo**

National ICT Australia  
School of CSE, UNSW  
Sydney, Australia  
tmeyer@cse.unsw.edu.au

**Rex Kwok**

Knowledge Systems Group  
School of CSE, UNSW  
Sydney, Australia  
rkwok@cse.unsw.edu.au

**Dongmo Zhang**

School of Computing and  
Information Technology  
University of Western Sydney, Australia  
dongmo@cit.uws.edu.au

## Abstract

This paper provides a logical framework for negotiation between agents that are assumed to be rational, cooperative and truthful. We present a characterisation of the permissible outcomes of a process of negotiation in terms of a set of rationality postulates, as well as a method for constructing exactly the rational outcomes. The framework is extended by describing two modes of negotiation from which an outcome can be reached. In the concessionary mode, agents are required to weaken their demands in order to accommodate the demands of others. In the adaptationist mode, agents are required to adapt to the demands of others in some appropriate fashion. Both concession and adaptation are characterised in terms of rationality postulates. We also provide methods for constructing exactly the rational concessions, as well as the rational adaptations. The central result of the paper is the observation that the outcomes obtained from the concessionary and adaptationist modes both correspond to the rational outcomes. We conclude by pointing out the links between negotiation and AGM belief change, and providing a glimpse of how this may be used to define a notion of preference-based negotiation.

## Introduction

Intelligent software agents involved in bargaining and negotiation on behalf of human clients are a reality (Sandholm 2002). As a result, negotiation is currently being investigated from many perspectives, including economics, applied mathematics, psychology, sociology and computer science (Bui & Shakun 1996; Rosenschein & Zlotkin 1994; Kraus, Sycara, & Evenchik 1998; Kraus 2001; Parsons, Sierra, & Jennings 1998). Thus far, most successful approaches have been quantitative in nature. Such approaches are usually game-theoretic in nature, with numeric utility functions forming the basis for decision-making. In many cases, however, numeric utilities are either unreliable or simply unavailable. This paper is a contribution to the body of literature, such as (Sycara 1990; Kraus, Sycara, & Evenchik 1998; Parsons, Sierra, & Jennings 1998; Booth 2001; 2002; Zhang *et al.* 2004) which, instead, views negotiation in a *qualitative* light. Our aim is to provide a logical framework for describing a process of negotiation between agents. The framework shares some similarities with (Wooldridge &

Parsons 2000) as well as with work on belief merging (Booth 2001; 2002). It is closely related to and, indeed, inspired by the work in (Zhang *et al.* 2004). For a more detailed treatment of negotiation see (Walton & Krabbe 1995).

In this paper we restrict ourselves to the case involving only two agents. Each agent comes to the negotiation table with an initial set of demands. The parties involved then proceed with a process of negotiation which terminates when they *strike a deal*—that is, when they have converged on a mutually agreed upon set of demands. The purpose of this paper is to define a logical framework for describing exactly those deals that are deemed to be rational.

Crudely speaking, the goal of an agent is two-pronged: it aims to have as many of its initial demands included in the negotiated outcome. At the same time it is also driven to reach an agreement acceptable to all parties. The framework we propose is intended to strike the correct balance between these two, possibly conflicting, goals. Agents are assumed to be truthful, rational and cooperative.

We commence with the provision of a set of rationality postulates, constraining the possible *outcomes* of the negotiation process between two agents. This is followed by a discussion of two different *modes* of negotiation, both also defined in terms of sets of rationality postulates. The first mode is a *concessionary* one. If the demands of two agents conflict, each is required to weaken its demands, by dropping some, in order to come to an agreement. The final negotiated outcome then consists of the *combination* of those demands that each agent chooses to retain. The second mode, termed *adaptation*, is one in which an agent is willing to deviate from its current demands in order to ensure that a settlement is reached. In the case of adaptation the negotiated outcome contains *only* those demands occurring in both adapted sets of demands. We show that the two modes of negotiation are interdefinable and, perhaps surprisingly, interdefinable with the framework for defining negotiated outcomes as well. We provide methods of constructing rational outcomes, concessions, and adaptations and prove, via appropriate representation results, that the sets of rationality postulates characterise these construction methods.

The results in this paper provide a basic logical framework for the definition of qualitative negotiation, but it does not address the question of defining the negotiation process itself, which is left as future work. In the conclusion we

elaborate on this, providing a brief discussion on the connection between negotiation and mutual belief revision, and how this connection may be employed to extend the currently defined framework to a preference-based one.

The logic under consideration is finitely generated and propositional, with the language denoted by  $\mathcal{L}$ , falsum by  $\perp$ , logical entailment by  $\models$ , logical equivalence by  $\equiv$  and logical closure by  $Cn$ . A *theory* is a set of sentences closed under logical entailment.  $M(K)$  denotes the *models* of a set of sentences  $K$  and  $M(\alpha)$  that of a single sentence  $\alpha$ .

A deal  $D$  is defined as an abstract object. Any deal is defined with respect to a *demand pair*  $K = (K_0, K_1)$ , with  $K_i$  ( $i = 0, 1$ ), being a *consistent* theory of  $\mathcal{L}$ , representing the initial demands, or demand set, of agent  $i$ . By doing so we are implicitly assuming that our agents are ideal reasoners, aware of all the logical consequences of their explicit demands. We add more structure to the formal definition of a deal as we proceed through the paper.

### Negotiated outcomes

We commence with the assumption that negotiating agents are ultimately interested in the *outcome* of the process of negotiation, irrespective of the mode of negotiation they adopt. Formally, the outcome  $O(D)$  of a deal  $D$  is a set of sentences representing the demands which both agents have agreed upon. A deal  $D$  is *outcome-permissible* iff  $O(D)$  satisfies the following rationality postulates:

- (O1)  $O(D) = Cn(O(D))$
- (O2)  $O(D) \not\models \perp$
- (O3) If  $K_0 \cup K_1 \not\models \perp$  then  $O(D) = Cn(K_0 \cup K_1)$
- (O4)  $(K_0 \cap K_1) \subseteq O(D)$  or  $O(D) \cup (K_0 \cap K_1) \models \perp$

(O1) ensures that outcomes are theories and (O2) that outcomes are consistent. (O3) states that if the initial demand sets do not conflict, an agent is obliged to uncritically add all demands posed by the other to its own, a requirement which might seem unrealistic. After all, why should I be forced to accept *all* demands of the other party involved just because our initial demands do not conflict? The response to this objection hinges on the current (admittedly simplistic) assumption that agents are cooperative. Other agents are seen as participants in a process in which the ultimate aim is to reach common ground. Cooperation implies the intention to meet as many of the demands of the other agent, as long as it does not conflict directly with one's own interests. (O4) prohibits the outcome of any process of negotiation to be consistent with the initial demands that both agents have in common without including *all* these commonly held demands. This can be justified by pointing out that if a potential outcome  $O$  is *consistent* with the demands that the agents have in common, it would be to the benefit of both rather to strengthen  $O$  to contain *all* commonly held demands.

### A classification of deals

The constraints placed on outcomes by (O1)-(O4) lead naturally to a taxonomy of deals in which we distinguish between four kinds of deals. A *trivial* deal is one for which the




Figure 1: Semantic representations of the outcome-permissible non-trivial deals.

outcome is  $Cn(K_0 \cup K_1)$ . This can only occur when the combination of the initial demand sets is consistent. In the next type of deal one of the agents, the *master*, gets to keep all its demands. An *i-dominant* deal ( $i = 0, 1$ ) is one in which agent  $i$  plays the role of the master. Observe that the different supersets of the demands of agent  $i$  correspond to the different  $i$ -dominant deals. The third type of deal is the class of *cooperative* deals, where the outcome is consistent with the initial demand set of each agent. Finally, we have the class of *neutral* deals, in which the outcome of negotiation is inconsistent with that which is common to the original demand sets. Figure 1 contains semantic representations of the non-trivial deals.

**Definition 1** Consider any deal  $D$ . If  $K_0 \cup K_1 \not\models \perp$  then  $D$  is the trivial deal, where  $O(D) = Cn(K_0 \cup K_1)$ . Otherwise we have the following classification.

1.  $D$  is an  $i$ -dominant deal iff  $O(D) \models K_i$ , for  $i = 0, 1$ .
2.  $D$  is a neutral deal iff  $O(D) \cup (K_0 \cap K_1) \models \perp$ .
3.  $D$  is a cooperative deal iff  $K_i \cup O(D) \not\models \perp$ , for  $i = 0, 1$ .

The semantic representation is particularly instructive in the case of the neutral deals. The intention is that a neutral deal will be struck when the starting positions of the two agents are too far apart for them to accommodate each other. In the spirit of cooperation they are then required to move to “neutral territory”, which is defined as any theory that is inconsistent with those initial demands that the agents have in common. It is in the justification of the choice of neutral territory that the semantic representation comes in useful: an outcome is neutral with respect to the initial demands of the two agents when it does not share any models with either of the initial demand sets.

It is easily established that the classification above provides a partition of the space of outcome-permissible deals.

**Theorem 1** Let  $D$  be any outcome-permissible deal. Then  $D$  is exactly one of a trivial,  $i$ -dominant ( $i = 0, 1$ ), neutral, or cooperative deal. Conversely, every deal which is trivial,  $i$ -dominant ( $i = 0, 1$ ), neutral, or cooperative is also outcome-permissible.

This classification is useful when setting protocols for negotiation, an issue that will be dealt with in future work.

## Constructing outcomes

We now describe a method for constructing those outcomes sanctioned by (O1)-(O4). Let  $\mathcal{OD} = \{D \mid O(D) \in \mathcal{O}\}$  where  $\mathcal{O} = \{Cn(K_0 \cup K_1)\}$  if  $K_0 \cup K_1 \not\perp \perp$ , and  $\mathcal{O} = \{\phi \mid \phi \not\perp \perp \text{ and } (Cn(\phi) \models K_0 \cap K_1 \text{ or } \phi \cup (K_0 \cap K_1) \models \perp)\}$ , otherwise.

As the next result shows, the set  $\mathcal{OD}$  contains precisely the outcome-permissible deals.

**Theorem 2** *A deal  $D$  satisfies (O1)-(O4) iff  $D \in \mathcal{OD}$ .*

Whenever the demand sets of two agents do not conflict, the postulates (O1)-(O4) thus allow the agents to converge on any consistent set either containing the commonly held demands, or in conflict with the commonly held demands.

**Example 1** *Let  $D$  be a deal such that  $K_0 = Cn(p \wedge q)$  and  $K_1 = Cn(\neg p \wedge \neg q)$ . Then  $O(D) \in \{Cn(p \wedge q), Cn(\neg p \wedge \neg q), Cn(p \wedge \neg q), Cn(\neg p \wedge q), Cn(p \leftrightarrow \neg q)\}$ . More specifically,  $D$  is an 0-dominated deal iff  $O(D) = Cn(p \wedge q)$ ,  $D$  is 1-dominated iff  $O(D) = Cn(\neg p \wedge \neg q)$ ,  $D$  is cooperative iff  $O(D) = Cn(p \leftrightarrow q)$  and  $D$  is neutral iff  $O(D) \in \{Cn(p \wedge \neg q), Cn(\neg p \wedge q), Cn(p \leftrightarrow \neg q)\}$ .*

Having laid down the basic requirements of negotiated outcomes, we now proceed with the description of two modes for obtaining such outcomes.

## Concession

In this section we propose a mode of negotiation which is *concessionary* in nature. Agents are expected to *weaken* their initial demands to reach an outcome acceptable to both parties. For a deal  $D$ , let  $C(D) = (C_0(D), C_1(D))$  be a pair of theories where  $C_i(D)$  ( $i = 0, 1$ ) represents the weakened demands, or *concessions*, of agent  $i$ , also referred to as an  $i$ -concession. The outcome of the negotiation is taken to be  $Cn(C_0(D) \cup C_1(D))$ . Formally, this means that we insist that the following property holds.

**(OC)**  $O(D) = Cn(C_0(D) \cup C_1(D))$

The link between outcomes and concessions will be strengthened further once the adaptationist mode of negotiation has also been discussed.

A deal  $D$  is *concession-permissible* iff  $C(D)$  satisfies the following rationality postulates:

**(C1)**  $C_i(D) = Cn(C_i(D))$  for  $i = 0, 1$

**(C2)**  $C_i(D) \subseteq K_i$  for  $i = 0, 1$

**(C3)** If  $K_0 \cup K_1 \not\perp \perp$  then  $C_i(D) = K_i$  for  $i = 0, 1$

**(C4)**  $C_0(D) \cup C_1(D) \not\perp \perp$

**(C5)** If  $C_0(D) \cup K_1 \not\perp \perp$  or  $C_1(D) \cup K_0 \not\perp \perp$  then  $K_0 \cap K_1 \subseteq C_0(D) \cup C_1(D)$

**(C6)** If  $C_0(D) \cup K_1 \models \perp$  and  $C_1(D) \cup K_0 \models \perp$  then  $C_0(D) \cup C_1(D) \cup (K_0 \cap K_1) \models \perp$

(C1) requires a concession to be a theory. (C2) ensures that the concession of an agent is a logical weakening of its initial demand set. (C3) deals with the case in which the initial demand sets are consistent. In this case neither agent has anything to gain by a weakening of its demand set. Note that (C3) expresses the same requirement as (O3) under the assumption that (OC) holds. (C4) is the insistence




Figure 2: A semantic representation of the concessions in  $CC$ . Observe that  $M(C_1(D)) = M(K_1)$  in the top left diagram, while  $M(C_0(D)) = M(K_0)$  in the top right diagram.

that the two agents reach some common ground when negotiating. It is the concessionary analogue of (O2), given that (OC) holds. (C5) is an appeal to the cooperative spirit in which negotiation is assumed to take place. Whenever one agent's concession is consistent with the other's initial demands, this postulate ensures that the demands they initially had in common are all included in the demands they have in common after both have conceded. Observe that in the presence of (C2) the entailment in the consequent of (C5) becomes equivalence. (C6) is a fairness requirement with a dichotomous flavour, and with rationality as the underlying motivation. If the concession of each agent is *inconsistent* with the initial demands of the other, then the final outcome (the two concessions combined) should be inconsistent with the initial demands they had in common. In other words, if neither of us were willing to weaken our position enough to become consistent with the initial demands of the other, we then agree to move to neutral territory.<sup>1</sup>

## Constructing concessions

We now consider a method for constructing the concession-permissible deals. Let  $\mathcal{CD} = \{D \mid C(D) \in \mathcal{C}\}$  where

$$\mathcal{C} = \begin{cases} \{(K_0, K_1)\}, & \text{if } K_0 \cup K_1 \not\perp \perp, \\ CC \cup IC & \text{otherwise} \end{cases}$$

where  $CC = \{(Cn(K_0 \cap \phi_1), Cn(K_1 \cap \phi_0)) \mid \phi_0 \not\perp \perp \text{ or } \phi_1 \not\perp \perp, \text{ and } \phi_i \models K_i \text{ for } i = 0, 1\}$

and  $IC = \{(Cn(K_0 \cap \phi), Cn(K_1 \cap \phi)) \mid \phi \not\perp \perp \text{ and } \phi \cup (K_0 \cap K_1) \models \perp\}$ .

As shown in Figure 2,  $CC$  contains exactly those cases in which both agents retain the initial commonly held beliefs, and at least one agent weakens its initial demand set so as not to conflict with the initial demands of the other. And as shown in Figure 3,  $IC$  contains those cases in which each agent picks the same set that conflicts with their commonly held demands, and restricts its initial demands to those occurring in this set as well.

<sup>1</sup>See the section on deal classification for a discussion of "neutral territory".




Figure 3: A semantic representation of concessions in  $IC$ .

**Example 2** Let  $D$  be such that  $K_0 = Cn(p \wedge q)$  and  $K_1 = Cn(\neg p \wedge \neg q)$ .  $CC = \{(Cn(p \wedge q), Cn(p \leftrightarrow q)), (Cn(p \leftrightarrow q), Cn(\neg p \wedge \neg q)), (Cn(p \leftrightarrow q), Cn(p \leftrightarrow q))\}$  and  $IC = \{(Cn(p), Cn(\neg q)), (Cn(q), Cn(\neg p)), (Cn(p \vee q), Cn(\neg p \vee \neg q))\}$ .

As the next result shows, the set  $CD$  contains precisely the concession-permissible deals.

**Theorem 3** A deal  $D$  satisfies (C1)-(C6) iff  $D \in CD$ .

We now have two classes of deals to consider—the outcome-permissible and concession-permissible deals. Before we look at the relationship between these, it is necessary to consider a second mode of negotiation.

### Adaptation

In this section we turn our attention to the *adaptationist* mode, where both agents are required to adapt their initial demands in such a way as to accommodate the demands of the other. For a deal  $D$ , let  $A(D) = (A_0(D), A_1(D))$  be a pair of theories where the adapted demands of agent  $i$  ( $i = 0, 1$ ), also referred to as an  $i$ -adaptation, are denoted by  $A_i(D)$ . The outcome of the negotiation process after adaptation has taken place consists of the sentences the two adaptation have in common. This is formalised as follows.

**(OA)**  $O(D) = A_0(D) \cap A_1(D)$

As we shall see,  $A_i(D)$  may conflict with  $K_i$ , thereby capturing the intuition that an agent's adaptive demands may differ radically from its initial demands. Such a situation is tenable only because of the assumption that both agents are cooperative and *know* that the other party is cooperative. An agent is therefore willing to risk a drastic departure from its initial demands towards the stated demands of the other only because it knows that the other party will not abuse its trust, and that both will therefore benefit from such a move.

A deal  $D$  is *adaptation-permissible* iff  $A(D)$  satisfies the following rationality postulates:

**(A1)**  $A_i(D) = Cn(A_i(D))$  for  $i = 0, 1$

**(A2)**  $K_0 \cup K_1 \not\equiv \perp$  implies  $A_0(D) = A_1(D) = Cn(K_0 \cup K_1)$

**(A3)**  $K_0 \subseteq A_i(D)$ , or  $K_1 \subseteq A_i(D)$ , or  $A_i(D) \cup (K_0 \cap K_1) \equiv \perp$ , for  $i = 0, 1$

**(A4)** For  $i = 0, 1$ , if  $K_i \not\subseteq A_i(D)$  then  $A_0(D) = A_1(D)$

(A1) ensures that adaptations are theories, while (A2) defines adaptation to be the combination of all initial demand




Figure 4: A semantic representation of the adaptations in  $CA$ .

in the case where the initial demands sets are not conflicting. Given (OA), it expresses the same requirement as (O3) and (C3). (A3) is a dichotomy principle based on rationality as motivation. It ensures that an agent's adaptation either includes at least one of the initial demand sets, or is inconsistent with the initial commonly held demands. The intuition is that an agent has exactly three choices when adapting: including all of its own demands, moving over completely to the viewpoint of the other agent, or otherwise moving to some neutral set of demands.<sup>2</sup> (A4) draws on the assumption of cooperation. If both agents do not include their own initial demands in their respective adaptations, they are required to cooperate completely by settling on identical adaptations.

Observe that the following property is a consequence of these postulates.

**(Rationality)**  $A_i(D) \models K_i \cap A_{1-i}(D)$  for  $i = 0, 1$

(Rationality) was proposed in (Zhang *et al.* 2004). It states that those initial demands of an agent adopted by its opponent should be retained in its own adaptation as well.

### Constructing adaptations

We now provide a method for constructing the adaptation-permissible deals. Let  $\mathcal{AD} = \{D \mid A(D) \in \mathcal{A}\}$  where

$$\mathcal{A} = \begin{cases} \{(Cn(K_0 \cup K_1), Cn(K_0 \cup K_1))\} & \text{if } K_0 \cup K_1 \not\equiv \perp, \\ C\mathcal{A} \cup I\mathcal{A} & \text{otherwise} \end{cases}$$

where  $C\mathcal{A} = \{(Cn(\phi_0), Cn(\phi_1)) \mid \phi_i \not\equiv \perp, \phi_i \models K_i, i = 0, 1\}$ , and  $I\mathcal{A} = \{(Cn(\phi), Cn(\phi)) \mid (\phi \cup (K_0 \cap K_1)) \equiv \perp, \phi \models K_0 \text{ or } \phi \models K_1, \text{ and } \phi \not\equiv \perp\}$

**Example 3** Let  $D$  be a deal such that  $K_0 = Cn(p \wedge q)$  and  $K_1 = Cn(\neg p \wedge \neg q)$ . Then  $C\mathcal{A} = \{(Cn(\neg p \wedge \neg q), Cn(p \wedge q))\}$ , and  $I\mathcal{A} = \{(Cn(p \wedge q), Cn(p \wedge q)), (Cn(\neg p \wedge \neg q), Cn(\neg p \wedge \neg q)), (Cn(p \wedge \neg q), Cn(p \wedge \neg q)), (Cn(\neg p \wedge q), Cn(\neg p \wedge q)), (Cn(p \leftrightarrow \neg q), Cn(p \leftrightarrow \neg q))\}$ .

Observe that  $C\mathcal{A}$  contains those cases in which each agent adopts a consistent superset of the other's initial demands. Figure 4 contains a semantic picture of such cases.

$I\mathcal{A}$ , on the other hand, contains those cases in which both agents adopt the *same* consistent set, provided it either contains the initial demands of one of the agents, or it conflicts

<sup>2</sup>See the section on deal classification for a discussion of "neutral territory".

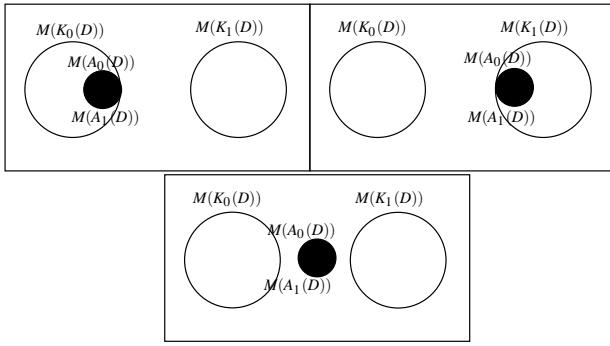


Figure 5: A semantic representation of the adaptations in  $\mathcal{IA}$ .

with the commonly held initial demands (a move to neutral territory). Figure 5 contains a semantic picture of such cases.

As the next result shows, the set  $\mathcal{AD}$  contains precisely the adaptation-permissible deals.

**Theorem 4** *A deal  $D$  satisfies (A1)-(A4) iff  $D \in \mathcal{AD}$ .*

### Outcome, concession and adaptation

We remarked earlier on our assumption that agents are primarily interested in the outcome of the negotiation process and that the mode of negotiation is of secondary concern. Nevertheless, the two different modes of negotiation are important in that they represent different approaches to the problem of characterising negotiation. Of course, measuring whether adaptation and concession yield the same results will have to be done via their common denominator, the *outcome* of the process of negotiation. Based on this, we make the reasonable assumption that both modes of negotiation should produce the same outcome, and it thus follows immediately from (OA) and (OC) that the next property holds:

$$(AC) \quad Cn(C_0(D) \cup C_1(D)) = A_0(D) \cap A_1(D)$$

(AC) insists that the combination of the concessions be equivalent to that which the adaptations have in common.

It is now possible to prove the following important result.

**Theorem 5**  *$O\mathcal{D} = C\mathcal{D} = \mathcal{AD}$ . That is, the classes of outcome-permissible deals, concession-permissible deals and adaptation-permissible deals are identical.*

In view of theorem 5 we shall refer to an element of any of these classes as *permissible*.

(AC) establishes a relationship between the concession and adaptation associated with a permissible deal, say  $D$ , but it is possible to go further and define  $C(D)$  and  $A(D)$  in terms of each other. This can be done with the following constructions.

$$(\text{Def C from A}) \quad C_i(D) = K_i \cap A_{1-i}(D), \text{ for } i = 0, 1$$

$$(\text{Def A from C}) \quad A_i(D) = \{Cn(C_i(D) \cup C_{1-i}(D)) \text{ if } K_i \cup C_{1-i}(D) \neq \perp, \text{ and } A_i(D) = Cn(K_i \cup C_{1-i}(D)) \text{ otherwise (for } i = 0, 1).$$

**Theorem 6** *For every permissible deal  $D$ ,  $C(D)$  can be defined in terms of  $A(D)$  and  $K(D)$  using (Def C from A), and  $A(D)$  can be defined in terms of  $C(D)$  and  $K(D)$  using (Def A from C).*

In fact, while we already know that an outcome  $O(D)$  obtained for a permissible deal  $D$  can be defined in terms of  $C(D)$  and in terms of  $A(D)$ , it is also possible to define  $C(D)$ , as well as  $A(D)$  in terms of the outcome  $O(D)$ . This can be done with the following constructions.

$$(\text{Def C from O}) \quad C_i(D) = K_i \cap O(D), \text{ for } i = 0, 1$$

$$(\text{Def A from O}) \quad A_i(D) = \{Cn(O(D) \cup K_i \text{ if } O(K) \cup K_i \neq \perp, \text{ and } A_i(D) = O(K) \text{ otherwise (for } i = 0, 1).$$

**Theorem 7** *For every permissible deal  $D$ ,  $C(D)$  can be obtained from  $O(D)$  and  $K(D)$  using (Def C from O), and  $A(D)$  can be obtained from  $O(D)$  and  $K(D)$  using (Def A from O).*

The next example illustrates the results of theorems 6 and 7.

**Example 4** *Consider any permissible deal  $D$  for which  $K_0 = Cn(p \wedge q)$  and  $K_1 = Cn(\neg p \wedge \neg q)$ . By theorem 5 and the results in examples 1, 2 and 3 we know that*

$$O(D) \in \left\{ \begin{array}{ll} Cn(p \wedge q), & Cn(\neg p \wedge \neg q), \\ Cn(p \leftrightarrow q), & Cn(p \wedge \neg q), \\ Cn(\neg p \wedge q), & Cn(p \leftrightarrow \neg q) \end{array} \right\},$$

$$C(D) \in \left\{ \begin{array}{l} (Cn(p \wedge q), Cn(p \leftrightarrow q)), \\ (Cn(p \leftrightarrow q), Cn(\neg p \wedge \neg q)), \\ (Cn(p \leftrightarrow q), Cn(p \leftrightarrow q)), \\ (Cn(p), Cn(\neg q)), \\ (Cn(q), Cn(\neg p)), \\ (Cn(p \vee q), Cn(\neg p \vee \neg q)) \end{array} \right\} \text{ and}$$

$$A(D) \in \left\{ \begin{array}{l} (Cn(p \wedge q), Cn(p \wedge q)), \\ (Cn(\neg p \wedge \neg q), Cn(\neg p \wedge \neg q)), \\ (Cn(\neg p \wedge \neg q), Cn(p \wedge q)), \\ (Cn(p \wedge \neg q), Cn(p \wedge \neg q)), \\ (Cn(\neg p \wedge q), Cn(\neg p \wedge q)), \\ (Cn(p \leftrightarrow \neg q), Cn(p \leftrightarrow \neg q)) \end{array} \right\}$$

By theorems 6 and 7 we know that

$$O\mathcal{D} = C\mathcal{D} = \mathcal{AD} = \{D_1, D_2, D_3, D_4, D_5, D_6\}$$

where:

$$O(D_1) = Cn(p \wedge q), \quad O(D_2) = Cn(\neg p \wedge \neg q), \\ O(D_3) = Cn(p \leftrightarrow q), \quad O(D_4) = Cn(p \wedge \neg q), \\ O(D_5) = Cn(\neg p \wedge q), \quad O(D_6) = Cn(p \leftrightarrow \neg q)$$

and

$$C(D_1) = (Cn(p \wedge q), Cn(p \leftrightarrow q)), \\ C(D_2) = (Cn(p \leftrightarrow q), Cn(\neg p \wedge \neg q)), \\ C(D_3) = (Cn(p \leftrightarrow q), Cn(p \leftrightarrow q)), \\ C(D_4) = (Cn(p), Cn(\neg q)), \\ C(D_5) = (Cn(q), Cn(\neg p)), \\ C(D_6) = (Cn(p \vee q), Cn(\neg p \vee \neg q))$$

and

$$A(D_1) = (Cn(p \wedge q), Cn(p \wedge q)), \\ A(D_2) = (Cn(\neg p \wedge \neg q), Cn(\neg p \wedge \neg q)), \\ A(D_3) = (Cn(\neg p \wedge \neg q), Cn(p \wedge q)), \\ A(D_4) = (Cn(p \wedge \neg q), Cn(p \wedge \neg q)), \\ A(D_5) = (Cn(\neg p \wedge q), Cn(\neg p \wedge q)), \\ A(D_6) = (Cn(p \leftrightarrow \neg q), Cn(p \leftrightarrow \neg q))$$

## Conclusion and future work

This paper has established a basic framework for determining permissible deals. We conclude with a brief discussion of how such a framework is to be used in a process of negotiation. The details will be the subject of future research.

The process of reaching an agreement involves a sequence of negotiation steps where, at each step, both agents simultaneously put forward the deal for which they currently have the highest preference. The process terminates when both agents are willing to accept the most recent deal proposed by the other party. Such a process assumes that, at any point, each agent has a preference relation on deals. Roughly speaking, these preferences are analogous to the utility functions used in quantitative settings. While the preferences of an agent can vary considerably, it has to take into account the demands and preferences of the other party involved in the negotiation process as well. The initial demands are known when the negotiation process commences, but the only other information available to an agent is the sequence of proposed deals it obtains from the other party. An agent thus has to make use of the information contained in these deals, and combine it with its knowledge about the negotiation process to perform a kind of preference extraction to guide it in its choice of deals to propose. This is one of the main topics to be dealt with in future work.

## Negotiation as belief revision

The reader familiar with the area of belief change will be struck by the similarities between adaptation and concession on the one hand, and AGM belief revision and contraction on the other (Alchourrón, Gärdenfors, & Makinson 1985; Gärdenfors 1988).<sup>3</sup> By conceding, an agent weakens its current demands, just as a belief contraction operation weakens the beliefs of an agent to some acceptable level. By adapting, an agent either strengthens its current demands, adopts a set which includes the demands of its adversary, or settles on a set that is inconsistent with the initial commonly held demands. Compare this with belief revision in which an agent either strengthens its own beliefs, or adopts a set which includes the sentence with which to revise and is inconsistent with its original beliefs.

These similarities are no coincidence. It can be shown that negotiation, as defined in this paper, can also be described as mutual belief revision (Meyer *et al.* 2004). Roughly speaking, the idea is the following. If the initial demands of the agents are conflicting, each agent is required to present a weakened version of their demands to the other which, in turn, is obliged to accept this weaker set of demands, by revising its current demands with it. It turns out that a suitably defined set of basic AGM revision operations can be used to obtain the adaptations of all permissible deals.

One of the advantages of looking at adaptation as basic AGM revision is that it provides a crude notion of preference among deals. This notion can be refined by a move

to those AGM revision operations satisfying the supplementary AGM postulates as well—a move that proves to be very useful in the definition of appropriate negotiation protocols.

## References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50:510–530.
- Booth, R. 2001. A negotiation-style framework for non-prioritised revision. In van Benthem, J., ed., *Theoretical Aspects of Rationality and Knowledge: Proceedings of the Eighth Conference (TARK 2001)*, 137–150. San Francisco, California: Morgan Kaufmann.
- Booth, R. 2002. Social contraction and belief negotiation. In *Proceedings of KR'02*, 375–384.
- Bui, T. X., and Shakun, M. F. 1996. Negotiation processes, evolutionary systems design, and negotiator. In Shakun, M. F., ed., *Negotiation Processes*. Kluwer Academic Publishers. 39–53.
- Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge, Massachusetts: The MIT Press.
- Kraus, S.; Sycara, K.; and Evenchik, A. 1998. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence* 104:1–69.
- Kraus, S. 2001. *Strategic Negotiation in Multiagent Systems*. MIT Press.
- Lau, R.; Essam, B.; Chan, S.; and Huang, Z. 2003. Belief revision for adaptive negotiation agents. 196–202.
- Meyer, T.; Foo, N.; Kwok, R.; and Zhang, D. 2004. Logical foundations of negotiation: strategies and preferences. In *Proceedings of KR'04*.
- Parsons, S.; Sierra, C.; and Jennings, N. 1998. Agents that reason and negotiate by arguing. *Journal of Logic and Computation* 8(3):261–292.
- Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter*. MIT Press.
- Sandholm, T. 2002. e-mediator: A next generation electronic commerce server. *Computational Intelligence* 18(4):656–667.
- Sycara, K. 1990. Persuasive argumentation in negotiation. *Theory and Decision* 28:203–242.
- van der Meyden, R. 1994. Mutual belief revision (preliminary report). In *Proceedings of KR'94*, 595–606.
- Walton, D., and Krabbe, E. C. W. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany, State University of New York Press.
- Wooldridge, M., and Parsons, S. 2000. Languages for negotiation. In *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)*, 393–400.
- Zhang, D.; Foo, N.; Meyer, T.; and Kwok, R. 2004. Negotiation as mutual belief revision. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-04)*.

<sup>3</sup>See (van der Meyden 1994) for work on mutual belief revision and (Lau *et al.* 2003) for connections between negotiation and belief revision.