

Logical Foundations of Negotiation: Strategies and Preferences

Thomas Meyer and Norman Foo

National ICT Australia, Sydney Node
School of Computer Science
and Engineering
University of New South Wales
Sydney, Australia
{tmeyer,norman}@cse.unsw.edu.au

Rex Kwok

Knowledge Systems Group
School of Computer Science
and Engineering
University of New South Wales
Sydney, Australia
rkwok@cse.unsw.edu.au

Dongmo Zhang

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

Abstract

This paper is a contribution towards the body of literature which views negotiation in a qualitative light. It builds on an existing logical framework for negotiation between rational, cooperative, truthful agents proposed in (Meyer, Kwok, & Zhang 2003). We show that agents equipped with negotiation strategies corresponding to basic AGM belief revision operations are capable of reaching exactly the *permissible* deals defined and discussed in (Meyer, Kwok, & Zhang 2003). Each agent has to present a set of weakened demands to the other party who, in return, is obliged to accept all weakened demands. The acceptance of demands is modelled by basic AGM belief revision.

We extend the logical framework of (Meyer, Kwok, & Zhang 2003) by considering scenarios in which the initial demand sets of agents may vary. We show that it forces agents to limit their negotiation strategies to AGM belief revision satisfying the supplementary AGM postulates. This leads to the redefinition of a negotiation strategy as a *preference relation* on demands. This extended framework provides a description of the deals that an agent ought to consider as reasonable, but provides no information on how it should go about choosing a particular deal. We conclude with suggestions on how negotiation strategies can be used to do so.

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 (Pruitt 1981; Bui & Shakun 1996; Rosenschein & Zlotkin 1994; Kraus, Sycara, & Evenchik 1998; Faratin, Sierra, & Jennings 1998; 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. 2003) which, instead, views negotiation in a *qualitative* light. It builds on the logical framework for qualitative negotiation in (Meyer, Kwok, & Zhang 2003) which, in turn, was inspired by the work in (Zhang *et al.* 2003). The framework shares some similarities with (Wooldridge & Parsons 2000) as well as with (Booth 2001; 2002). For a more detailed treatment of negotiation see (Walton & Krabbe 1995). Our purpose is to provide a framework in which agents attempt to reach agreement on a set of demands.¹ Agents are assumed to be truthful, rational and cooperative. Agents are *truthful* in the sense that the demands they pose at any time during the negotiation process are exactly those demands that they truly want to see fulfilled, and a deal they propose is the deal they truly deem to be the most preferred at that stage. Agents operate under a *principle of rationality*: they attempt to maximise their own gains without being concerned about whether other agents obtain ‘more’ out of the process than they do. And agents obey a *principle of cooperation* which requires of them to accommodate the demands of others, provided that it does not conflict with theirs.

We take the view that an agent is equipped with a *negotiation strategy* to aid in its quest for striking a deal. One of the central results of this paper is that such a negotiation strategy should be an AGM belief revision operation (Alchourrón, Gärdenfors, & Makinson 1985; Gärdenfors 1988). We show that in doing so, all and only the deals deemed to be *permissible* in (Meyer, Kwok, & Zhang 2003) will be generated.

We extend the logical framework of (Meyer, Kwok, & Zhang 2003) by considering scenarios in which the initial demands of agents may vary. To do so we consider entities referred to as *compound deals*. A compound deal contains deals corresponding to every possible pair of initial demand sets that two agents may have. In the second main result of the paper we show that the constraints we place on compound deals force agents to limit their negotiation strategies to AGM belief revision satisfying the supplementary AGM postulates. This leads to the redefinition of a negotiation strategy as a *preference relation* on demands.

The framework proposed for negotiation provides a description of the deals that an agent ought to consider as reasonable, but provides no information on how it should go

¹The consequences of such an agreement might include the actual allocation of resources.

about choosing a deal. We conclude with suggestions on how the negotiation strategies, redefined as preferences relations, can be employed to do so.

After dispensing with some formal preliminaries we review the logical framework for negotiation of (Meyer, Kwok, & Zhang 2003) on which our work is based. This is followed by a review of AGM belief change (Gärdenfors 1988). In the next section we recast basic AGM belief revision operations as negotiation strategies and show that they can be used to generate precisely the permissible deals presented in (Meyer, Kwok, & Zhang 2003). These results are then extended by showing that full AGM belief revision operations (those satisfying the supplementary postulates as well) can be used to generate a strict subclass of the permissible deal satisfying some additional constraints. This is followed by a discussion of related work presented in (Booth 2001; 2002). Finally, the last section concludes and considers future research.

Formal preliminaries

We take the logic under consideration to be a finitely generated propositional logic with \perp being the canonical representation of inconsistency. We denote logical entailment by \models , logical equivalence by \equiv , and logical closure by Cn . A set of sentences closed under logical entailment is referred to as a *theory*. The *models* of a set of sentences K is denoted by $M(K)$.

A logical framework for negotiation

In this section we briefly review the logical framework for negotiation presented in (Meyer, Kwok, & Zhang 2003). We consider a situation in which two rational, cooperative, truthful agents commence with a process of negotiation which terminates when they *strike a deal*. Both aim to have as many of their initial demands included in the final negotiated outcome, but are also driven to reach an agreement acceptable to all parties as quickly as possible.

The reader is referred to (Meyer, Kwok, & Zhang 2003) for details concerning motivation and justification. Associated with any deal is, firstly, the initial demands of the negotiators. But a deal also has an outcome associated with it – the common set of demands that both parties have agreed upon. Furthermore, we shall consider two modes of obtaining the outcome of a deal. In the first mode, participants are required to concede by retracting some of their initial demands. The outcome of a deal is then obtained by combining the demands that remain after both agents have made some concessions. Accordingly, we shall associate with every deal the concessions used to arrive at the outcome for that deal. In the second mode, agents have to adapt their demands in some appropriate fashion to reach an agreement. Unlike the concessionary mode, adaptation does not involve the retraction of initial demands. The adaptation process is driven by the knowledge that the final outcome consists of those demands common to the adaptation of the two agents and the assumption of cooperation between agents. We shall associate with every deal the adaptations used to obtain the outcome of a deal as well.

Formally, a deal D is then defined as an abstract object with respect to a *demand pair* $K = (K_0, K_1)$, with K_i ($i = 0, 1$), a theory, representing the initial demands of agent i . By doing so we assume that our agents are ideal reasoners, aware of all the logical consequences of their explicit demands.

The process of negotiation has an *outcome* $O(D)$, a set of sentences representing the demands which both agents have agreed upon. A deal D is said to be *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 $O(D) \cup (K_0 \cap K_1) \models \perp$

Outcomes have to be consistent theories. If the initial demands of the two agents are consistent, the outcome is obtained by combining all demands. If not, the outcome should either contain those demands common to the two initial demand sets, or should be inconsistent with the demands common to the initial demand sets.

The *concessions* associated with deals are formalised as follows. The concession $C(D) = (C_0(D), C_1(D))$ associated with a deal D is a pair of theories where $C_i(D)$ ($i = 0, 1$), referred to as an *i -concession*, represents the weakened demands of agent i . The outcome of the process of negotiation is obtained by combining the concessions of the two agents:

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

The permissible deals can also be characterised in terms of concessions. If we insist that (OC) holds, the permissible deals are exactly those satisfying 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 $K_0 \cup K_1 \not\models \perp$ then $C_i(D) = K_i$ for $i = 0, 1$
- (C4) $C_0(D) \cup C_1(D) \not\models \perp$
- (C5) If $C_0(D) \cup K_1 \not\models \perp$ or $C_1(D) \cup K_0 \not\models \perp$ then $K_0 \cap K_1 \subseteq C_0(D) \cap C_1(D)$
- (C6) If $C_0(D) \cup K_1 \models \perp$ and $C_1(D) \cup K_0 \models \perp$ then $C_0(D) \cup C_1(D) \cup (K_0 \cap K_1) \models \perp$

The concession of an agent has to be a theory weaker than its original demands. If the two original demands sets are consistent with each other, neither agent needs to concede at all. If not, the two *concessions* need to be consistent with each other. If the concession of one agent is consistent with the initial demands of the other, the demands common to the initial two demands sets should be included in the final outcome: the combination of the two concessions. If the concession of each agent is inconsistent with the initial demands of the other, the final outcome should be inconsistent with the demands common to the initial demand sets.

The *adaptations* associated with deals are formalised as follows. The adaptation $A(D) = (A_0(D), A_1(D))$ associated with a deal D is a pair of theories where $A_i(D)$

($i = 0, 1$), referred to as an i -adaptation, represents the adapted demands of agent i . This process is motivated by the requirement that the outcome consists of the demands the two adaptations have in common:

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

The permissible deals can also be characterised in terms of their adaptations. If we insist that (OA) holds, then the permissible deals are exactly those satisfying the following rationality postulates:

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

(A2) If $K_0 \cup K_1 \not\equiv \perp$ then

$$A_0(D) = A_1(D) = Cn(K_0 \cup K_1)$$

(A3) $K_0 \subseteq A_i(D)$, or $K_1 \subseteq A_i(D)$, or

$$A_i(D) \cup (K_0 \cap K_1) \equiv \perp, \text{ for } i = 0, 1$$

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

Adaptations have to be theories. If the initial demand sets do not conflict, an adaptation is obtained by aggregating all the initial demands. An adaptation includes one of the two initial demands sets, or is inconsistent with the demands common to the initial demand sets. Finally, if the adaptation of an agent does not include its own initial demands, the adaptations of both agents have to be identical.

Note that the adoption of both (OA) and (OC) means that the combined concessions (and their consequences) should be exactly those sentences that the adaptations of the agents have in common:

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

The next important step is to provide a classification of the permissible deals. A *trivial deal* D is one for which the outcome is $Cn(K_0 \cup K_1)$. This occurs when $Cn(K_0 \cup K_1)$ is consistent. If $Cn(K_0 \cup K_1)$ is inconsistent, the permissible deals are partitioned into the following categories.

1. An i -dominated deal D ($i \in \{0, 1\}$) is one in which the outcome $O(D)$ includes the initial demands K_i of agent i . Clearly agent i dominates in such a deal.
2. A *cooperative deal* D is one in which the outcome is consistent with the initial demands of both agents (i.e. $O(D)$ is consistent with K_0 and consistent with K_1).
3. A *neutral deal* D is one in which the outcome is inconsistent with the initial demands of both agents (i.e. $O(D)$ is inconsistent with K_0 and inconsistent with K_1).

It is easily established that the classification above provides a partition of the space of permissible deals. Figure 1 contains semantic representations of the outcomes of the permissible non-trivial deals.²

AGM belief change

AGM belief change (Alchourrón, Gärdenfors, & Makinson 1985; Gärdenfors 1988) investigates the rational ways for an agent to change its mind. The current set of beliefs of an agent is represented by a theory K . AGM belief change

²See (Meyer, Kwok, & Zhang 2003) for semantic representations of the concessions and adaptation of permissible deals.

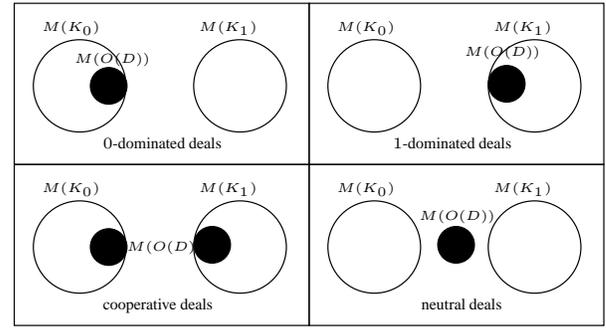


Figure 1: Semantic representations of the outcomes of the permissible non-trivial deals.

is concerned with two forms of belief change: revision, in which an agent has to incorporate new information while maintaining consistency, and contraction, in which an agent has to remove information from its current beliefs. The reader is referred to (Gärdenfors 1988) for more details.

As pointed out by Nayak (Nayak 1994), AGM belief revision can also be thought of as the revision of a theory by a theory. This is due to the inclusion of the properties of *closure* and the *irrelevance of syntax*. For the rest of the paper we shall adopt this approach and consider the revision of a theory K by a theory C . Given our assumption of a finitely generated propositional logic, the following characterisation of *basic* AGM belief revision is equivalent to the original formulation in (Alchourrón, Gärdenfors, & Makinson 1985), where K is the theory to be revised, $*$ is the revision operation, and C the input with which to revise the theory K .

(K*1) $K * C = Cn(K * C)$

(K*2) $K * C \subseteq Cn(K \cup C)$

(K*3) If $K \cup C \not\equiv \perp$ then $K * C = Cn(K \cup C)$

(K*4) $C \subseteq K * C$

(K*5) $C = Cn(C)$

(K*6) $K * C \equiv \perp$ iff $C \equiv \perp$

Basic AGM belief revision accepts theories as input and produces consistent theories, except in the pathological case where the input is the inconsistent theory. It always includes the input in the resulting theory, which means it is always successful. If the input C is consistent with the original set of beliefs, the revised theory is obtained by simply adding the elements of C to K and closing under logical consequence.

The postulates for basic AGM revision are all concerned with a fixed input C . *Full* AGM revision can be made more systematic by insisting on the following *supplementary* postulates.

(K*7) $K * Cn(A \cup B) \subseteq Cn((K * A) \cup B)$

(K*8) If $(K * A) \cup B \not\equiv \perp$ then

$$Cn((K * A) \cup B) \subseteq K * Cn(A \cup B)$$

One of the great benefits of including the supplementary postulates is an important result linking up full AGM revision with preferences on valuations. Let \preceq be a total preorder on valuations and denote the \preceq -minimal valuations of a theory K by $M_{\preceq}(K)$. Then \preceq is termed K -faithful iff the \preceq -minimal valuations are exactly the models of K (i.e. $M(K) = M_{\preceq}(Cn(\top))$). We then get the following representation result showing that the revision of a theory C is obtained by taking the theory generated by the \preceq -minimal models of C (see e.g. (Gärdenfors 1988; Katsuno & Mendelzon 1991)).

Theorem 1 *For every K -faithful total preorder on valuations \preceq , there is a full AGM revision operation $*$ such that $K * C = Th(M_{\preceq}(C))$ for every theory C . Conversely, for every full AGM revision operation $*$, there is a K -faithful total preorder on valuations \preceq such that $K * C = Th(M_{\preceq}(C))$ for every theory C .*

This result allows us to view every full AGM revision operation as a representation of the preferences of an agent with regard to its beliefs. It is well-known that the standard representation of preferences on beliefs in the belief revision literature—epistemic entrenchments (Gärdenfors 1988)—corresponds exactly to the faithful total preorders on valuations.

Negotiation strategies

The observant reader will be struck by the similarities between adaptation and concession on the one hand, and AGM belief revision and contraction on the other. By conceding, an agent weakens its current demands to some acceptable level, just as the contraction by a theory C weakens the beliefs of an agent to the level of not containing C . 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 the revision by a theory C in which an agent either strengthens its own beliefs, or adopts a set which includes C , and is inconsistent with its originally held beliefs. These similarities are not a coincidence. In this section we show that negotiation can be modelled as mutual belief revision (van der Meyden 1994).³ Roughly speaking, the idea is that whenever the initial demands of the agents are conflicting, each agent will be required to present a weakened version of their demands to the other which, in turn, is obliged to accept this weaker set of demands. It is the process of *accepting* weakened demands that will be modelled by AGM belief revision. In this context we shall refer to the AGM belief revision operations as the possible *negotiation strategies* of the agents. This process is feasible only because of the assumption of cooperation between agents. Agents will only be willing to accept the weakened demands of others unconditionally if they are assured that these demands are presented in good faith. And in turn, they will be obliged to reciprocate.

³See also (Lau et al. 2003).

Rational negotiation strategies

Our analysis commences with an investigation of whether *all* basic AGM revision operations should qualify as negotiation strategies. So let us assume that agent i has decided on \otimes as a negotiation strategy. If agent $i - 1$ now presents to i a set of weakened demands C which i has to accept, how will agent i revise its current demands? That is, what should $K_i \otimes C$ look like? Observe firstly that if C is consistent with K_i , then *every* basic AGM revision operation will produce $Cn(K_i \cup C)$ as the result. Hence, this case does not rule out any basic AGM revision operation. If the preferred demands of agent $(1 - i)$ are inconsistent with the demands of agent i (i.e. $K_i \cup C \models \perp$), the situation can be subdivided into those cases for which C is

1. equal to K_{1-i} (and inconsistent with K_i), or
2. strictly weaker than K_{1-i} (but inconsistent with K_i).

In the former case, agent $(1 - i)$ regards all demands in K_{1-i} as equally preferable, and this places no restriction on the choice of \otimes by agent i . In the latter case, agent $(1 - i)$ expresses a preference for the demands in C over the remaining demands in K_{1-i} . The principle of cooperation then requires of agent i to respect these preferences. So, the demands that $K_i \otimes C$ and K_{1-i} have in common should be exactly those found in C . That is, we require that

$$(A) \quad (K_i \otimes C) \cap K_{1-i} = C.$$

On the other hand, the principle of rationality requires the inclusion of as many demands as possible. Consequently $K_i \otimes C$ should be the *largest* set of sentences subject to the restriction imposed on it in (A).⁴

Example 1 *Let $K_0 = Cn(p \wedge q)$ and $K_1 = Cn(\neg p \wedge \neg q)$. Clearly $Cn(\neg q)$ is logically weaker than K_1 and is inconsistent with K_0 . Furthermore, the result of revising K_0 by $Cn(\neg q)$, using a basic AGM revision operation may produce any, and only, one of the following theories: $Cn(\neg q)$, $Cn(p \wedge \neg q)$, or $Cn(\neg p \wedge \neg q)$. Of these, only the first two satisfy (A). And of these, the largest is $Cn(p \wedge \neg q)$.*

In summary then, the principles of cooperation and rationality lead us to exclude some basic AGM revision operations from being considered as viable negotiation strategies. The remaining ones will be referred to as *rational*.

Definition 1 *The negotiation strategy \otimes for agent i ($i = 0, 1$) with demand set K_i is rational iff for every input C such that $K_i \cup C \models \perp$ and $C \subset K_{1-i}$, it is the case that $K_i \otimes C$ is the largest set satisfying (A).*

Permissible and compatible inputs

The next step is to consider what the valid *inputs* to a negotiation strategy might look like. The intuition is that the inputs to the negotiation strategy of agent i will correspond to the possible concessions of agent $(1 - i)$: the $(1 - i)$ -concession. If we allow the initial demand set of agent $(1 - i)$ to vary, any theory C will qualify as a valid input to the negotiation strategy. We shall consider such cases in the next section. But for a fixed demand pair $K = (K_0, K_1)$, certain inputs

⁴It is easy to establish the existence of such a set.

will be irrelevant since they will never be posed. For example, it is easy to check that for $K_0 = Cn(p \wedge q)$ and $K_1 = Cn(\neg p \wedge \neg q)$, agent 1 will never present $Cn(\top)$, or for that matter, $Cn(p \vee \neg q)$ or $Cn(\neg p \vee q)$, to agent 0 as a concession. The reason is simply that there is no permissible deal D for which $C_1(D)$ takes on any of these values.

Definition 2 A theory C is an i -permissible input ($i = 0, 1$) for a negotiation strategy iff there is a permissible deal D such that $C_{1-i}(D) = C$.

The negotiation strategy of an agent dictates how it should respond when presented with a set of (weakened) demands from its adversary, and it is thus a partial encoding of the preferences of the agent. But note that this same agent will also need to present such a set of weakened demands to the other party. If we want the agent to be rational, it also needs to take these preferences into account when choosing demands to present to its opponent. Consider an example where $K_0 = Cn(p \wedge q)$ and $K_1 = Cn(\neg p)$, and let \otimes be the (unique) rational negotiation strategy for which $K_0 \otimes Cn(\neg p) = Cn(\neg p \wedge q)$. That is, \otimes prefers $Cn(\neg p \wedge q)$ over all the other supersets of K_1 . Now, if agent 0 is rational, the same *kind* of preferences should still hold if it restricts each of these supersets to those demands occurring in K_0 . That is, agent 0 should prefer $Cn(\neg p \wedge q) \cap K_0$ over $Cn(\neg p) \cap K_0$ and $Cn(\neg p \wedge \neg q) \cap K_0$. It turns out that these are exactly the 1-permissible inputs (or 0-concessions) that are consistent with K_1 . In fact, this remark can be generalised as follows.

Proposition 1 Let $K = (K_0, K_1)$ be any demand pair and let $i \in \{0, 1\}$. For every $A \supseteq K_{1-i}$, $K_i \cap A$ is an $(1-i)$ -permissible input and is consistent with K_{1-i} . Conversely every $(1-i)$ -permissible input that is consistent with K_{1-i} , is equal to $K_{1-i} \cap A$ for some $A \supseteq K_{1-i}$.

What proposition 1 tells us then, is that of all i -concessions of permissible deals that are consistent with K_{1-i} , an agent i with \otimes as negotiation strategy will prefer the one that is equal to $K_i \cap (K_i \otimes K_{1-i})$. This leads to the following definition.

Definition 3 Let $K = (K_0, K_1)$ be any demand pair, $i \in \{0, 1\}$, and \otimes a rational negotiation strategy for agent i . The i -concession $C_i(D)$ of a permissible deal D is i -compatible with \otimes iff $C_i(D) \cup K_{1-i} \models \perp$ or $C_i(D) = K_i \otimes K_{1-i}$.

Determination of deals

Now that we have a clear picture of what constitutes a rational negotiation strategy and an input compatible with it, we aim to provide a proper correspondence between the rational negotiation strategies and the permissible deals. To do so, it remains to define the notion of a negotiation strategy *determining* a permissible deal. Intuitively the permissible deals determined by a negotiation strategy \otimes are the ones that an agent adopting \otimes will consider.

Definition 4 Let $K = (K_0, K_1)$ be a demand pair and $i \in \{0, 1\}$. A rational negotiation strategy \otimes i -determines a permissible deal D iff $K_i \otimes C_{1-i}(D) = A_i(D)$ and $C_i(D)$ is i -compatible with \otimes . Two permissible deals D and D' are i -codetermined by \otimes iff they are both i -determined by it. A

set of permissible deals \mathcal{D} is i -codetermined by \otimes iff the elements of \mathcal{D} are pairwise i -codetermined by \otimes . A permissible deal D is uniquely i -determined by \otimes iff it is i -determined by \otimes and i -codetermined only by itself.

The i -determination of a permissible deal D refers to the ability of agent i to generate $A_i(D)$ when presented with $C_{1-i}(D)$ as input, while simultaneously sticking to the i -concessions that are i -compatible with \otimes . Observe that codetermination partitions the set of permissible deals into equivalence classes, with each of the uniquely determined deals contained in its own partition.

It is important to realise that i -determination is subjective in the sense that it only provides us with agent i 's perspective on the situation. That is, the agent uses $K_i \otimes C_{1-i}(D)$ to determine $A_i(D)$, but \otimes does not have anything explicitly to say about $A_{1-i}(D)$, and may therefore not contain enough information to specify a unique permissible deal. This is why the notions of unique determination and codetermination are needed. If a permissible deal D is uniquely i -determined then, by definition, \otimes specifies D , and only D , when presented with the appropriate input. Contrast this with the case of two distinct permissible deals, say D and D' , that are i -codetermined by \otimes . By definition this means that \otimes is unable to distinguish between D and D' , and that some additional information would be needed to do so. The following result outlines when such additional information is needed.

Proposition 2 If $K_0 \cup K_1 \not\models \perp$ then every rational negotiation strategy uniquely i -determines the trivial deal, and only the trivial deal. Now suppose that $K_0 \cup K_1 \models \perp$, and consider any rational negotiation strategy \otimes . Then

1. A deal D is i -determined by \otimes iff $C_i(D)$ is i -compatible with \otimes .
2. \otimes uniquely i -determines every neutral deal.
3. \otimes uniquely i -determines exactly one $(1-i)$ -dominated deal D and does not i -determine any other $(1-i)$ -dominated deal.
4. \otimes i -determines every i -dominated deal.
5. For every i -dominated deal D there is a single cooperative deal D' for which $C_i(D')$ is i -compatible with \otimes , such that D and D' are i -codetermined by \otimes .
6. For every cooperative deal D for which $C_i(D)$ is i -compatible with \otimes , there is an i -dominated deal D' such that D and D' are i -codetermined by \otimes .
7. No two i -dominated deals are i -codetermined by \otimes .

The next example aims to illustrate the various notions relating to i -determination and provides concrete instances of the results in proposition 2.

Example 2 Let $K_0 = Cn(p \wedge q)$ and $K_1 = Cn(\neg p)$. The 0-permissible inputs are

$Cn(\neg p)$, $Cn(\neg p \vee \neg q)$, and $Cn(\neg p \vee q)$.

There is a single neutral deal D_N with

$O(D_N) = Cn(p \wedge \neg q)$,

one 0-dominated deal D_0 with

$O(D_0) = Cn(p \wedge q)$,

three 1-dominated deals D_1 , D_2 , and D_3 , with

$$\begin{aligned} O(D_1) &= Cn(\neg p \wedge q), \\ O(D_2) &= Cn(\neg p \wedge \neg q), \text{ and} \\ O(D_3) &= Cn(\neg p), \end{aligned}$$

and three cooperative deals D_{C1} , D_{C2} and D_{C3} , with

$$\begin{aligned} O(D_{C1}) &= Cn(q), \\ O(D_{C2}) &= Cn(p \leftrightarrow q), \text{ and} \\ O(D_{C3}) &= Cn(\neg p \vee q). \end{aligned}$$

Now, let \otimes be the (unique) rational negotiation strategy for which $K_0 \otimes Cn(\neg p) = Cn(\neg p \wedge q)$. The 0-concessions (or 1-permissible inputs), are

$$Cn(p \wedge q), Cn(p), Cn(q), Cn(p \leftrightarrow q), \text{ and } Cn(\neg p \vee q),$$

but only the first three are 0-compatible with \otimes .

It is easily verified that \otimes uniquely 0-determines the only neutral deal as well as the 1-dominated deal D_1 . Observe, however, that \otimes does not 0-determine the remaining two 1-dominated deals since

$$(K_0 \otimes C_1(D_2)) \neq A_0(D_2)$$

and

$$(K_0 \otimes C_1(D_3)) \neq A_0(D_3).$$

Furthermore, \otimes 0-determines the only 0-dominated deal, as well as the cooperative deal D_{C1} . In fact, these two deals are 0-codetermined by \otimes . To verify this claim, consider the following description of these two deals.

$$\begin{aligned} D_0: \quad & C_0(D_0) = Cn(p \wedge q), \\ & C_1(D_0) = Cn(\neg p \vee q), \\ & O(D_0) = A_0(D_0) = A_1(D_0) = K_0 \otimes C_1(D_0) \text{ where} \\ & K_0 \otimes C_1(D_0) = Cn(p \wedge q) \\ D_{C1}: \quad & C_0(D_{C1}) = Cn(q), \\ & C_1(D_{C1}) = Cn(\neg p \vee q), \\ & O(D_{C1}) = Cn(q), \\ & A_0(D_{C1}) = K_0 \otimes C_1(D_{C1}) = Cn(p \wedge q), \\ & A_1(D_{C1}) = Cn(\neg p \wedge q) \end{aligned}$$

Observe that $C_0(D_0)$ and $C_0(D_{C1})$ are both 0-compatible with \otimes . Furthermore, note that

$$A_0(D_0) = A_0(D_{C1}) = Cn(p \wedge q),$$

and so these two deals are not uniquely 0-determined, but are both 0-codetermined by \otimes .

Finally, neither of the two remaining cooperative deals, D_{C1} and D_{C2} , are 0-determined by \otimes , since neither $C_0(D_{C2})$ nor $C_0(D_{C3})$ are 0-compatible with \otimes . This can be verified by considering the description of these two deals below.

$$\begin{aligned} D_{C2}: \quad & C_0(D_{C2}) = Cn(p \leftrightarrow q), \\ & C_1(D_{C2}) = Cn(\neg p \vee q), \\ & O(D_{C2}) = Cn(p \leftrightarrow q), \\ & A_0(D_{C2}) = K_0 \otimes C_1(D_{C2}) = Cn(p \wedge q), \\ & A_1(D_{C2}) = Cn(\neg p \wedge \neg q) \\ D_{C3}: \quad & C_0(D_{C3}) = Cn(\neg p \vee q), \\ & C_1(D_{C3}) = Cn(\neg p \vee q), \\ & O(D_{C3}) = Cn(\neg p \vee q), \\ & A_0(D_{C3}) = K_0 \otimes C_1(D_{C3}) = Cn(p \wedge q), \\ & A_1(D_{C3}) = Cn(\neg p) \end{aligned}$$

From proposition 2 it thus follows that a rational negotiation strategy for agent i contains all the information needed to obtain unique neutral and $(1 - i)$ -dominated deals, but may not contain enough information to obtain unique cooperative and i -dominated deals. This is because, although \otimes cuts down on the i -concessions that agent i may present to agent $(1 - i)$, it may not contain enough information to specify a unique i -concession. In the case of neutral and $(1 - i)$ -dominated deals this does not matter, because a unique i -concession can be inferred from the other available information. But for the i -dominated and cooperative deals this is not the case. For example, consider the two deals D_0 and D_{C1} that are codetermined by \otimes in example 2 above. There is nothing in \otimes that indicates whether agent 0 should choose the 0-concession $C_0(D_0) = Cn(p \wedge q)$ or the 0-concession $C_0(D_{C1}) = Cn(q)$.

Proposition 2 places us in a position to formalise the connection between permissible deals and rational negotiation strategies.

Theorem 2.1. *For every pair of rational negotiation strategies (\otimes_0, \otimes_1) , there is a permissible deal that is i -determined by \otimes_i , for $i = 0, 1$.*

2. *For every permissible deal D there is a pair of rational negotiation strategies (\otimes_0, \otimes_1) such that D is i -determined by \otimes_i , for $i = 0, 1$.*

Part 1 of theorem 2 is a ‘‘soundness’’ result. It shows that all pairs of rational negotiation strategies will generate permissible deals. Part 2 is a ‘‘completeness’’ result. It shows that there is a way to generate every permissible deal using rational negotiation strategies. As a consequence, the rational negotiation strategies can be said to *characterise* the permissible deals.

Negotiation strategies and compound deals

The permissible deals are defined with respect to a fixed demand pair, but one would expect a systematic rational agent to take its original choice of deals into account when choosing deals with respect to *other* demand pairs as well. Let \mathcal{K} be an enumeration of all demand pairs $K = (K_0, K_1)$. We define a *compound deal* \mathcal{D} as an n -tuple, where n is the number of demand pairs, and for every $j \in \{1, \dots, n\}$, entry j in \mathcal{D} , denoted by D_j , is a permissible deal with respect to the j th demand pair $K^j = (K_0^j, K_1^j)$ in the enumeration \mathcal{K} . A compound deal represents the permissible deals corresponding to different demand pairs that an agent judges to be related in some way. In this section we investigate which compound deals a rational agent should accept as reasonable. A compound deal \mathcal{D} is called *permissible* iff the outcomes, concessions and adaptations of the deals in \mathcal{D} , all satisfy the following postulates:

(O5) $\forall j, k \in \{1, \dots, n\}, i \in \{0, 1\}$, if
 $K_i^j = K_i^k, K_{1-i}^j \subseteq K_{1-i}^k, K_{1-i}^j \subseteq O(D_j)$,
and $O(D_j) \cup K_{1-i}^k \neq \perp$,
then $O(D_k) = Cn(O(D_j) \cup K_{1-i}^k)$

(C7) $\forall j, k \in \{1, \dots, n\}, i \in \{0, 1\}$, if
 $K_i^j = K_i^k, K_{1-i}^j \subseteq K_{1-i}^k, C_i(D_j) \cup K_{1-i}^j \not\subseteq \perp$,
and $C_i(D_j) \cup K_{1-i}^k \not\subseteq \perp$,
then $C_i(D_k) = Cn(C_i(D_j) \cup K_{1-i}^k) \cap K_i^j$

(A5) $\forall j, k \in \{1, \dots, n\}, i \in \{0, 1\}$, if
 $K_i^j = K_i^k, K_{1-i}^j \subseteq K_{1-i}^k, K_{1-i}^j \subseteq A_i(D_j)$,
and $A_i(D_j) \cup K_{1-i}^k \not\subseteq \perp$,
then $A_i(D_k) = Cn(A_i(D_j) \cup K_{1-i}^k)$

It is perhaps best to motivate these postulates by thinking of $K^k = (K_0^k, K_1^k)$ as a “new” and modified version of the “old” demand pair $K^j = (K_0^j, K_1^j)$. An analysis of the postulates shows that all three are concerned only with $(1-i)$ -dominated deals, and with cases where the demands of agent i are unchanged, but where agent $(1-i)$ becomes more demanding (i.e. K_{1-i}^k in the new demand pair is logically stronger than K_{1-i}^j in the old demand pair). (O5) then requires that the new outcome be equal to the old outcome combined with the new stronger demands of agent $(1-i)$. Similarly, (A5) requires that the new i -adaptation be equal to the old i -adaptation combined with the new stronger demands of agent $(1-i)$. And finally, (C7) requires that a demand be admitted to the new i -concession iff it occurred in the old demand set of agent i , and it occurs in the combination of the old i -concession with the new stronger demands of agent $(1-i)$.

The significance of these postulates is that they enable us to link up the permissible compound deals with full AGM revision. Let us refer to those rational negotiation strategies that are full AGM revision operations, as *systematic* negotiation strategies. We generalise the idea of a negotiation strategy determining a deal as follows:

Definition 5 For $i \in \{0, 1\}$, a negotiation strategy \otimes i -determines a permissible compound deal \mathcal{D} iff for every $j \in \{1, \dots, n\}$, $K_i^j \otimes C_{1-i}(D_j) = A_i(D_j)$ and $C_i(D_j)$ is i -compatible with \otimes .

When a compound deal \mathcal{D} is i -determined it means that for every demand pair K^j , the negotiation strategy \otimes i -determines the deal D_j in \mathcal{D} . The notions of codetermination and unique determination extend to compound deals in the obvious way, and will not be defined explicitly here. Also, we shall not provide a result for compound deals that is analogous to proposition 2 at this time. But we are in a position to state the result showing that the systematic negotiation strategies *characterise* the permissible compound deals.

Theorem 3.1. For every pair of systematic negotiation strategies (\otimes_0, \otimes_1) , there is a permissible compound deal that is i -determined by \otimes_i , for $i = 0, 1$.

2. For every permissible compound deal \mathcal{D} there is a pair of systematic negotiation strategies (\otimes_0, \otimes_1) such that \mathcal{D} is i -determined by \otimes_i , for $i = 0, 1$.

By theorems 1 and 3 it follows that each systematic negotiation strategy can be redefined as a total preorder on valuations. As such, it can be viewed as an encoding of the

preferences of the agent who chooses to adopt it. But observe that since every systematic negotiation strategy is also a rational negotiation strategy, the class of full AGM belief revision operations available to an agent is constrained by all the factors discussed in the previous section. In this way the previous section helps to pick out those AGM revision operations particularly suited to negotiation.

Related work

In (Booth 2001; 2002) Richard Booth considers negotiation as a version of the problem of merging the inputs obtained from different sources. The agents involved are required to weaken their initial demands to the point where the resultant demands are jointly consistent. The outcome of the negotiation process is then taken as the closure of the combination of the weakened sets of demands. Booth shows that the process of weakening demands can be seen as a generalised version of AGM contraction, which he terms *social contraction*, and argues that combining the weakened demands results in a version of the Levi Identity (Levi 1991); a standard construction for obtaining a belief revision operation from a belief contraction operation. The weakened demands of agents in this setting are clearly analogous to the concessions of permissible deals. The concessions associated with permissible deals are all instances of Booth’s social contraction, although the converse does not hold. The main reason for this can be traced back to the postulates (C5) and (C6). These two postulates require of the set, say X , of demands the two agents originally have in common to exhibit dichotomous behaviour. Either X should be included in the demands commonly held after weakening has taken place (and therefore included in the final outcome of the deal), or it should be *inconsistent* with the final outcome. It turns out that social contraction does not always satisfy this requirement.

(Booth 2002) is concerned with a postulational version of this setting,⁵ while (Booth 2001) describes a procedural version in which agents progressively weaken their initially held demands until the combination is jointly consistent. So it is only at the termination of the process of reaching agreement that an outcome is encountered that satisfies the relevant postulates. In fact, termination of the process occurs *precisely* when such an outcome is found. This is in contrast to the scenario we envisage (and briefly discuss in the next section) in which an acceptable deal is reached through a process of navigating the space of permissible deals. One of the advantages of the kind of process we have in mind is that it lends itself to a game-theoretic analysis (Luce & Raiffa 1989), thereby providing access to the tools available in that discipline for measuring the quality of decisions.

Conclusion

The main results in this paper concern the use of basic AGM belief revision operations, and more importantly, *full* AGM belief revision operations as suitable negotiation strategies

⁵It also considers the general case of a finite number of negotiating agents.

to be used by an agent intent on specifying which deals it regards as reasonable. The result concerning full AGM belief revision is useful because it allows us to consider a negotiation strategy as a total preorder on valuations. And from results in classical AGM belief change (Gärdenfors 1988), this means that a negotiation strategy can be recast as an *epistemic entrenchment*, which can be used as a preference relation on demands.

While these results *describe* the deals that a rational agent would regard as reasonable, they currently do not provide information on how an agent should *choose* a particular deal. It is our contention that the view of a negotiation strategy as a preference relation on demands provides the basis for such a choice. One way in which to use such preference relations would be to define a process in which both agents repeatedly remove their least entrenched demands until they reach a state where the remaining two sets of demands are jointly consistent. This is, essentially, the procedure described in (Booth 2001), although the process there is not guided by an entrenchment relation. However, we intend to employ the preference relation in a different manner. Roughly speaking, the preferences on demands can be *lifted* in appropriate ways to preferences on *permissible deals*. The preferences on deals are then used as the basis for defining negotiation protocols. So, instead of an agent focusing on relaxing its own initial demands, such a process requires of agents to reason about and navigate the space of permissible deals. It is at this point where game-theoretic notions such as strategies dominating or being dominated, versions of Pareto optimality, and versions of equilibrium will come into play. However, the development and analysis of these negotiation protocols is left as future work.

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.
- Faratin, P.; Sierra, C.; and Jennings, N. R. 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24:159–182.
- Gärdenfors, P. 1988. *Knowledge in Flux : Modeling the Dynamics of Epistemic States*. Cambridge, Massachusetts: The MIT Press.
- Katsuno, H., and Mendelzon, A. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.
- Kraus, S.; Sycara, K.; and Evenchik, A. 1998. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence* 104:1–69.
- Lau, R.; Essam, B.; Chan, S.; and Huang, Z. 2003. Belief revision for adaptive negotiation agents. 196–202.
- Levi, I. 1991. *The Fixation of Belief and Its Undoing*. Cambridge: Cambridge University Press.
- Luce, R. D., and Raiffa, H. 1989. *Games and Decisions: Introduction and Critical Survey*. Dover Publications.
- Meyer, T.; Kwok, R.; and Zhang, D. 2003. Logical foundations of negotiation: Outcome, concession and adaptation. Norman Foo's Festschrift, <http://www.cse.unsw.edu.au/~ksg/Norman/>
- Nayak, A. C. 1994. Iterated belief change based on epistemic entrenchment. *Erkenntnis* 41:353–390.
- Parsons, S.; Sierra, C.; and Jennings, N. 1998. Agents that reason and negotiate by arguing. *Journal of Logic and Computation* 8(3):261–292.
- Pruitt, D. G. 1981. *Negotiation Behaviour*. Academic Press.
- 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. 2003. Negotiation as mutual belief revision. In *Proceedings of the Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC'03)*.