

A negotiation-style framework for non-prioritised revision

Richard Booth
University of Leipzig
Department of Computer Science
Augustusplatz 10/11
04109 Leipzig, Germany
booth@informatik.uni-leipzig.de

Abstract

We present a framework for non-prioritised belief revision — i.e., belief revision in which newly acquired information is not always fully accepted — in which the result of revision is arrived at via a kind of negotiation between old information and new. We show how both ordinary partial meet revision and Fermé and Hansson’s selective revision can be captured in this framework, and also how it supports the definition of contraction operators which do not necessarily satisfy the basic AGM contraction postulate of (Success).

1 Introduction

Belief revision has been an active area of research, both in philosophy and computer science, since the early 1980’s. The most popular basic framework within which it is studied is the one laid down by Alchourrón, Gärdenfors and Makinson (AGM) in [1]. One of the basic assumptions of that framework is that newly acquired information should always be accepted after revision and that any resulting inconsistency should be repaired by deleting some of the previous beliefs. In other words the new information is given “priority” over the old. This paper is inspired by recent developments in the area of so-called *non-prioritised* belief revision, i.e., belief revision in which newly acquired information is *not* always fully accepted (see [9] for a recent survey).

The basic problem of belief revision is the following: How should an agent change his beliefs — represented by a consistent, deductively closed set of sentences (belief set) K — in response to the new information that some consistent sentence ϕ is true, in such a way that the resulting set of beliefs is again consistent and deductively closed?¹ If ϕ is consistent with K then we can simply add

¹We are using the AGM representation of belief states and belief inputs, with the simpli-

ϕ to K and then close under logical consequence, but what if ϕ is inconsistent with K ? In this case we can imagine this inconsistency as a “gap” between K on one side and ϕ on the other. The problem of revision may then be thought of as the problem of how to bridge this gap, i.e., reach consistency. In partial meet revision, i.e., the revision operation of AGM theory, the bridge is constructed entirely from the side of K , in that enough of K is removed so that it becomes possible to consistently add ϕ (cf. the Levi Identity). This ensures ϕ is always fully accepted in the new belief set. In Fermé and Hansson’s *selective revision* [3], which itself generalises several approaches to non-prioritised revision (see [9]), a part of the bridge is first constructed from the side of ϕ , in that ϕ is transformed into some logically weaker sentence ϕ' (the “selected part” of ϕ , which is allowed to be logically equivalent to ϕ itself), and *then* the bridge is completed from the side of K , in that enough of K is removed so that ϕ' can be consistently added. In this type of revision ϕ' , but possibly *not* ϕ , will be accepted in the new belief set. The basic idea behind this paper is to push this bridge-building analogy further and consider an even more general model of revision which allows for the bridge to be constructed from both sides simultaneously.

To help us do this we first introduce an abstract framework for merging together pieces of information coming from two different sources. In this framework the merging itself can be seen as a kind of *negotiation* between the sources which resolves any inconsistency. The main construct here is that of a *belief negotiation model*. A given belief negotiation model, relative to the two information sources s and t , uniquely determines the course of any negotiation between s and t and thus uniquely determines the result of merging any two pieces of information provided by s and t respectively. By interpreting s as the belief-holding agent and t as an information source external to the agent, this framework becomes a framework for investigating non-prioritised revision which, as we will see, provides both a useful starting point from which different ideas on non-prioritised revision can be worked out, and an interesting perspective on the process of belief change itself.

The plan of this paper is as follows. In Section 2 we introduce our general merging framework and define belief negotiation models. Then in Section 3 we apply this framework to belief change, showing how a given belief negotiation model \mathcal{N} not only gives rise to a non-prioritised revision operator $\boxplus_{\mathcal{N}}$, but also yields an operator $\boxminus_{\mathcal{N}}$ of the other main type of belief-change operator studied in the AGM framework, i.e., *contraction*. Given a belief negotiation model \mathcal{N} we will call the pair $(\boxplus_{\mathcal{N}}, \boxminus_{\mathcal{N}})$ a *basic revision-contraction pair*. We give an axiomatic characterisation of such pairs in Section 3, where we also look at a few of their other properties, including how $\boxplus_{\mathcal{N}}$ and $\boxminus_{\mathcal{N}}$ interrelate. As we will see, $\boxplus_{\mathcal{N}}$ satisfies all the so-called basic AGM postulates for revision except, possibly, (Success), while $\boxminus_{\mathcal{N}}$ satisfies all the basic AGM postulates for contraction except, possibly, (Success) and (Recovery). Then in Section 4

ifying assumptions that the agent’s belief set is consistent, that he is never directed to revise his belief set by an inconsistent sentence, and, later, that he is never directed to contract his belief set by a tautology.

we add three restrictions in turn to belief negotiation models which lead to three progressively smaller classes of basic revision-contraction pairs. We give axiomatic characterisations for each of the three classes. The revision operators $\boxplus_{\mathcal{N}}$ of the smallest class correspond to partial meet revision operators while those of the middle class correspond to one particular class of the selective revision operators studied in [3]. The ones of the largest class, which satisfy a very weak form of the revision postulate (Success), have not appeared elsewhere. In Section 5 we briefly mention some related work from the literature on merging before concluding and giving a couple of ideas for further work in Section 6.

Before we get started let us fix our notation. We work throughout in a propositional language L generated from finitely many, but at least two,² propositional variables. Cn denotes the classical logical consequence operator on L . The (finite) set of all propositional worlds is denoted by \mathcal{W} and the set of all non-empty subsets of \mathcal{W} is denoted by \mathcal{B} . For any set $I \subseteq L$, $[I]$ denotes the set of worlds in which every sentence in I is true. If $[I] \neq \emptyset$ then I is consistent, otherwise I is inconsistent. For $\phi \in L$ we write $[\phi]$ rather than $[\{\phi\}]$ and write “ ϕ is consistent” rather than “ $\{\phi\}$ is consistent” etc. If $\neg\phi$ is inconsistent, i.e., $[\phi] = \mathcal{W}$, then ϕ is a tautology. A belief set is any set of sentences $K \subseteq L$ which is both consistent and deductively closed, i.e., $K = Cn(K)$. The set of all belief sets we denote by \mathcal{K} . We denote the set of all consistent sentences by L^* and the set of all non-tautologous sentences by L_* . (So $\phi \in L^*$ iff $\neg\phi \in L_*$.) Thus for us a revision operator will be a function with domain $\mathcal{K} \times L^*$ while a contraction operator will be a function with domain $\mathcal{K} \times L_*$. Given any $I \cup \{\phi\} \subseteq L$, $I + \phi$ denotes $Cn(I \cup \{\phi\})$ while, for any set of worlds $S \subseteq \mathcal{W}$, $Th(S)$ denotes the set of sentences true in every world in S . Finally, for any set X , $|X|$ denotes the cardinality of X .

2 Belief negotiation models

We begin by describing our basic merging framework.³ Suppose we have two sources of information s and t which impart the information S and T respectively, where S and T are subsets of \mathcal{W} to be interpreted as the information that the actual “true” world is one of the worlds in S , respectively T . How can we *merge* these two pieces of information into a single piece $\text{Merge}(S, T)$? This basic problem has already been considered, outside of the literature on non-prioritised revision, in several works (see Section 5). Our idea to solve it is to incrementally weaken, i.e., enlarge, S or T or both until “common ground” is reached, i.e., until their intersection is non-empty. To be slightly more precise we start off with the pair $\langle S_0, T_0 \rangle$ where $S_0 = S$ and $T_0 = T$. If $S_0 \cap T_0 \neq \emptyset$ then we just take $\text{Merge}(S, T) = S_0 \cap T_0$. But if $S_0 \cap T_0 = \emptyset$, i.e., S_0 and T_0 contradict each other, then we perform what may be thought of as a “round of negotiation” which is just a contest between S_0 and T_0 . The loser of this contest (or both

²Some of our later counter-examples seem to depend on this assumption.

³We remark that this framework shares some similarities with the abstract formalisation of negotiation found in [18]. For a more detailed treatment of negotiation see [17].

parties, if the contest ends in a draw) must then “make some concessions”, i.e., make some weakening of his position by admitting more possibilities, while the winner stays the same. Thus we arrive at the pair $\langle S_1, T_1 \rangle$ where $S_0 \subseteq S_1$ and $T_0 \subseteq T_1$. Now if $S_1 \cap T_1 \neq \emptyset$ then we set $\text{Merge}(S, T) = S_1 \cap T_1$. Otherwise the next round of negotiation takes place, this time between S_1 and T_1 . Once again the loser of this round makes concessions, and we keep going like this until $S_i \cap T_i \neq \emptyset$, at which point we set $\text{Merge}(S, T) = S_i \cap T_i$. Let us spell all this out in detail.

We assume that the items of information supplied by s and t are always non-empty subsets of \mathcal{W} , i.e., elements of \mathcal{B} . Let Ω denote the set of all finite sequences of pairs of elements of \mathcal{B} . Given $\omega = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle) \in \Omega$ we will say ω is *increasing* iff $S_i \subseteq S_{i+1}$ and $T_i \subseteq T_{i+1}$ for all $i = 0, 1, \dots, n-1$. We define the set of sequences $\Sigma \subseteq \Omega$ by

$$\Sigma = \{ \omega = (\langle S_0, T_0 \rangle, \langle S_1, T_1 \rangle, \dots, \langle S_n, T_n \rangle) \in \Omega \mid \omega \text{ is increasing,} \\ \text{and } S_n \cap T_n = \emptyset \}.$$

A sequence $\sigma = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle) \in \Sigma$ represents a possible stage in the unfinished (since $S_n \cap T_n = \emptyset$) negotiation process starting with S_0 and T_0 . Given $i < n$, we let σ_i denote that sequence consisting of the first $i+1$ entries in σ , i.e., $\sigma_i = (\langle S_0, T_0 \rangle, \dots, \langle S_i, T_i \rangle)$.

In the simple negotiation scenario described above there were two ingredients in need of further specification. Firstly, we need to know how a round of negotiation is carried out. In this paper we will only be interested in the *results* of such negotiation rounds. Thus for our purposes it suffices to assume the existence of a function $g : \Sigma \rightarrow 2^{\mathcal{B}}$ which satisfies, for each $\sigma \in \Sigma$,

$$(g0) \quad \emptyset \neq g(\sigma) \subseteq \{S_n, T_n\}, \text{ where } \sigma = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle).$$

The intuition behind the function g is that it selects, at each negotiation stage σ , which of the two parties should make concessions. In other words it returns the loser of the negotiation round at stage σ (or both parties, if the round ends in a draw). In order to avoid deadlock, we stipulate that at least one party must make concessions, i.e., $g(\sigma) \neq \emptyset$.

The second problem is then to decide what concessions the loser of the negotiation round should make. Once again we abstract away from the actual process used to determine this and assume only that we are given, for each $\sigma = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle) \in \Sigma$, a function $\nabla_\sigma : \{S_n, T_n\} \rightarrow \mathcal{B}$ with the interpretation that $\nabla_\sigma(S_n)$ represents the weakening of S_n given that S_n must be weakened at stage σ , and similarly for $\nabla_\sigma(T_n)$. Once again to avoid deadlock, we stipulate that this weakening is strict, i.e., that ∇_σ satisfies the following condition for each $A \in \{S_n, T_n\}$:

$$(\nabla 0) \quad A \subseteq \nabla_\sigma(A) \text{ and } \nabla_\sigma(A) \not\subseteq A^4$$

We can now make the following definition.

⁴Note that, even though we are requiring that $\nabla_\sigma(A)$ be a strict weakening of A for both $A = S_n$ and $A = T_n$, these weakenings will only actually be “carried out” if A loses the

Definition 1 A belief negotiation model (relative to s and t) is a pair $\mathcal{N} = \langle g, \{\nabla_\sigma\}_{\sigma \in \Sigma} \rangle$ where $g : \Sigma \rightarrow 2^{\mathcal{B}}$ is a function which satisfies (g0) and, for each $\sigma = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle) \in \Sigma$, $\nabla_\sigma : \{S_n, T_n\} \rightarrow \mathcal{B}$ is a function which satisfies ($\nabla 0$).

In what follows we will sometimes denote the function g of a belief negotiation model $\mathcal{N} = \langle g, \{\nabla_\sigma\}_{\sigma \in \Sigma} \rangle$ by $g_{\mathcal{N}}$.

A belief negotiation model \mathcal{N} then uniquely determines, for any given $S, T \in \mathcal{B}$ provided by s and t respectively, the complete process of negotiation between S and T . This process is returned by the function $f^{\mathcal{N}} : \mathcal{B} \times \mathcal{B} \rightarrow \Omega$ given by

$$f^{\mathcal{N}}(S, T) = \sigma = (\langle S_0, T_0 \rangle, \langle S_1, T_1 \rangle, \dots, \langle S_n, T_n \rangle)$$

where (i) $S_0 = S$ and $T_0 = T$, (ii) n is minimal such that $S_n \cap T_n \neq \emptyset$, and (iii) for each $0 \leq i < n$ we have

$$S_{i+1} = \begin{cases} \nabla_{\sigma_i}(S_i) & \text{if } S_i \in g(\sigma_i) \\ S_i & \text{otherwise.} \end{cases} \quad \text{and} \quad T_{i+1} = \begin{cases} \nabla_{\sigma_i}(T_i) & \text{if } T_i \in g(\sigma_i) \\ T_i & \text{otherwise.} \end{cases}$$

A belief negotiation model thus gives us a way of merging two items of information S and T . We simply take $\text{Merge}(S, T) = S_n \cap T_n$. However an interesting aspect of our approach, which we will exploit in the next section, is that it allows us to study S_n and T_n in their own right. For this purpose we define, from a given belief negotiation model \mathcal{N} , the functions $f_{\rightarrow}^{\mathcal{N}}, f_{\leftarrow}^{\mathcal{N}} : \mathcal{B} \times \mathcal{B} \rightarrow \mathcal{B}$ by

$$f_{\rightarrow}^{\mathcal{N}}(S, T) = S_n \quad \text{and} \quad f_{\leftarrow}^{\mathcal{N}}(S, T) = T_n,$$

where $f^{\mathcal{N}}(S, T) = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle)$. Thus $f_{\rightarrow}^{\mathcal{N}}(S, T)$ is the result (according to \mathcal{N}) of weakening information S from source s to accommodate information T from source t , while $f_{\leftarrow}^{\mathcal{N}}(S, T)$ is the result of weakening information T from source t to accommodate information S from source s . Returning to our bridge analogy from the introduction, $f_{\rightarrow}^{\mathcal{N}}(S, T)$ represents that part of the bridge between S and T constructed from the side of S while $f_{\leftarrow}^{\mathcal{N}}(S, T)$ is the part built from the side of T . Let us try and illustrate these definitions with an example.

Example 1 Let p and q be distinct propositional variables in L . Let's assume we know source s to be more reliable than source t when it comes to information concerning the truth value of p , while s and t are equally reliable concerning the truth value of q . Now suppose s provides the information that p and q are both false, i.e., $S = [\neg p \wedge \neg q]$ while t provides the conflicting information that p and q are both true, i.e., $T = [p \wedge q]$. We can merge these two pieces of information using a belief negotiation model $\mathcal{N} = \langle g, \{\nabla_\sigma\}_{\sigma \in \Sigma} \rangle$ relative to s and t by constructing the sequence $f^{\mathcal{N}}(S, T) = \sigma$ in stages as follows. Let $S_0 = S$, $T_0 = T$. Then $\sigma_0 = (\langle S_0, T_0 \rangle)$. Since $S_0 \cap T_0 = \emptyset$ some negotiation

negotiation round at stage σ , i.e., $A \in g(\sigma)$. Hence to avoid deadlock it is really only necessary that $\nabla_\sigma(A)$ be a strict weakening of A for some $A \in g(\sigma)$. Our stronger requirement above comes from our desire to keep our conditions on the ∇_σ independent from our conditions on g .

is necessary. We assume the first round of negotiation proceeds by focusing on the truth value of p , about which there is obvious disagreement between S_0 and T_0 . Since we know information about p provided by s is more reliable than such information provided by t we conclude that T_0 is more likely to be wrong here and so T_0 is declared the loser of this negotiation round. This is reflected in \mathcal{N} by setting $g(\sigma_0) = \{T_0\}$. At this stage T_0 is weakened by giving up p , but holding on to q , i.e., $\nabla_{\sigma_0}(T_0) = [q] = T_1$, while $S_0 = S_1$ remains unchanged. Thus we reach the stage $\sigma_1 = (\langle S_0, T_0 \rangle, \langle S_1, T_1 \rangle)$. Since $S_1 \cap T_1 = \emptyset$, a second round of negotiation is required in which, this time, we focus on the truth value of q . Again there is disagreement here between S_1 and T_1 , but this time, since s and t are equally reliable in matters concerning q , no outright winner can be called in this round. This is reflected by setting $g(\sigma_1) = \{S_1, T_1\}$. Hence both S_1 and T_1 must be weakened. We assume S_1 gives up $\neg q$ but holds on to $\neg p$, i.e., $\nabla_{\sigma_1}(S_1) = [\neg p] = S_2$, while T_1 gives up q but holds on to $p \rightarrow q$, i.e., $\nabla_{\sigma_1}(T_1) = [p \rightarrow q] = T_2$. Now we have $S_2 \cap T_2 \neq \emptyset$ and so the negotiation process ends here. We thus have

$$f^{\mathcal{N}}([\neg p \wedge \neg q], [p \wedge q]) = (\langle [\neg p \wedge \neg q], [p \wedge q] \rangle, \langle [\neg p \wedge \neg q], [q] \rangle, \langle [\neg p], [p \rightarrow q] \rangle),$$

while $f^{\mathcal{N}}([\neg p \wedge \neg q], [p \wedge q]) = [\neg p]$ and $f^{\mathcal{N}}([\neg p \wedge \neg q], [p \wedge q]) = [p \rightarrow q]$. The result of merging $[\neg p \wedge \neg q]$ and $[p \wedge q]$, according to \mathcal{N} , is given by

$$f^{\mathcal{N}}([\neg p \wedge \neg q], [p \wedge q]) \cap f^{\mathcal{N}}([\neg p \wedge \neg q], [p \wedge q]) = [\neg p].$$

Note that all values of g and ∇_{σ} for $\sigma \in \Sigma$, other than the ones mentioned, are irrelevant to this example.

We now turn to applying belief negotiation models in a belief change setting.

3 Basic revision-contraction pairs

One way to view non-prioritised revision is as an operation which merges old information (i.e., that contained in the agent's belief set) with new information. This suggests performing revision using a belief negotiation model \mathcal{N} relative to the information sources s and t , where now s is the agent and t is the agent's external information source. Formally we define the revision operator $\boxplus_{\mathcal{N}}$ from \mathcal{N} by, for $K \in \mathcal{K}$ and $\phi \in L^*$,

$$K \boxplus_{\mathcal{N}} \phi = Th(f^{\mathcal{N}}_{\rightarrow}([K], [\phi]) \cap f^{\mathcal{N}}_{\leftarrow}([K], [\phi])).$$

But note that a spin-off of using belief negotiation models in this context is that they also give rise to a notion of *contraction*. Interpreting the contraction of K by $\phi \in L^*$ as the weakening of K to accommodate $\neg\phi$ leads us quite naturally to define, from \mathcal{N} , the contraction operator $\boxminus_{\mathcal{N}}$ by

$$K \boxminus_{\mathcal{N}} \phi = Th(f^{\mathcal{N}}_{\rightarrow}([K], [\neg\phi])).$$

We make the following definition.

Definition 2 Let $\boxplus : \mathcal{K} \times L^* \rightarrow 2^L$ and $\boxminus : \mathcal{K} \times L_* \rightarrow 2^L$ be a pair of functions. Then (\boxplus, \boxminus) is a basic revision-contraction pair iff $\boxplus = \boxplus_{\mathcal{N}}$ and $\boxminus = \boxminus_{\mathcal{N}}$ for some belief negotiation model \mathcal{N} .

In the AGM theory of belief change, operators of revision \boxplus and contraction \boxminus are each assumed to satisfy several properties known as the *basic AGM postulates for revision and contraction*⁵ which we shall list now. We refer to each postulate by its usual name, with the addition that we will prefix revision postulates by r- and contraction postulates by c-. The basic AGM postulates for revision are:

- $K \boxplus \phi = Cn(K \boxplus \phi)$ (r-Closure)
- If $\phi_1 \leftrightarrow \phi_2 \in Cn(\emptyset)$ then $K \boxplus \phi_1 = K \boxplus \phi_2$ (r-Extensionality)
- $\phi \in K \boxplus \phi$ (r-Success)
- $K \boxplus \phi \subseteq K + \phi$ (r-Inclusion)
- If $K \cup \{\phi\}$ is consistent then $K + \phi \subseteq K \boxplus \phi$ (r-Vacuity)
- $K \boxplus \phi$ is consistent (r-Consistency)

Any function $\boxplus : \mathcal{K} \times L^* \rightarrow 2^L$ satisfying the above postulates is known as a *partial meet revision operator* [1].⁶ For contraction the basic postulates are:

- $K \boxminus \phi = Cn(K \boxminus \phi)$ (c-Closure)
- If $\phi_1 \leftrightarrow \phi_2 \in Cn(\emptyset)$ then $K \boxminus \phi_1 = K \boxminus \phi_2$ (c-Extensionality)
- $\phi \notin K \boxminus \phi$ (c-Success)
- $K \boxminus \phi \subseteq K$ (c-Inclusion)
- If $\phi \notin K$ then $K \boxminus \phi = K$ (c-Vacuity)
- $K \subseteq (K \boxminus \phi) + \phi$ (c-Recovery)

Any function $\boxminus : \mathcal{K} \times L_* \rightarrow 2^L$ satisfying the above contraction postulates is known as a *partial meet contraction operator* [1].

For operators of non-prioritised revision, of course, the postulate (r-Success) is considered too strong. Similarly for contraction one can imagine situations in which (c-Success) is not satisfied, for example if an agent holds a belief so firmly that it becomes “irretractible” from his belief set [16].⁷ Another contraction

⁵We do not consider the so-called *supplementary* postulates in this paper. See [8] for a description of these as well as an explanation of the basic postulates.

⁶Actually the usual formulation of (r-Consistency) is “if ϕ is consistent then $K \boxplus \phi$ is consistent” but, since we do not allow revising by inconsistent sentences, the antecedent here is vacuous for us. Similar remarks apply to the contraction postulate (c-Success), whose usual formulation is “if ϕ is not a tautology then $\phi \notin K \boxminus \phi$ ”.

⁷We remark that study of operations of contraction which violate (c-Success) does not seem to have received the same level of attention as its revision counterpart. Exceptions here are [4] and, using the framework of default logic, [5].

postulate which has been open to much debate is (c-Recovery) (see e.g. [6, 13]). Any function $\boxminus : \mathcal{K} \times L_* \rightarrow 2^L$ which satisfies all the basic AGM contraction postulates with the possible exception of (c-Recovery) is known as a *withdrawal operator* [13].

Another characteristic of AGM theory is the interdefinability of revision and contraction. To define a revision operator from a contraction operator we may use the Levi Identity, while for the other direction we have the Harper Identity:

$$K \boxplus \phi = (K \boxminus \neg\phi) + \phi \quad (\text{Levi Identity})$$

$$K \boxminus \phi = (K \boxplus \neg\phi) \cap K \quad (\text{Harper Identity})$$

What can we say about the properties of a basic revision-contraction pair? We have the following syntactic characterisation.

Theorem 1 *Let $\boxplus : \mathcal{K} \times L^* \rightarrow 2^L$ and $\boxminus : \mathcal{K} \times L_* \rightarrow 2^L$ be a pair of functions. Then the following are equivalent:*

- (i). $\langle \boxplus, \boxminus \rangle$ is a basic revision-contraction pair.
- (ii). \boxplus satisfies (r-Closure), (r-Extensionality), (r-Inclusion), (r-Vacuity) and (r-Consistency); \boxminus satisfies (c-Closure), (c-Extensionality), (c-Inclusion) and (c-Vacuity); and $\langle \boxplus, \boxminus \rangle$ satisfies, for all $K \in \mathcal{K}$ and $\phi \in L^*$,

$$K \boxplus \phi \subseteq (K \boxminus \neg\phi) + \phi \quad (1/2\text{-Levi})$$

$$K \boxminus \neg\phi \subseteq K \boxplus \phi \quad (\text{Mixed Inclusion})$$

The reader will note that the only basic AGM postulates missing from part (ii) in the above theorem are the two success postulates and (c-Recovery).

In the next section we will see what happens to $\boxplus_{\mathcal{N}}$ and $\boxminus_{\mathcal{N}}$ when we force certain restrictions on \mathcal{N} . Before we do that, however, let's look a little more closely at the behaviour of basic revision-contraction pairs.

3.1 Other properties

In the rest of this section, unless otherwise indicated, $\langle \boxplus, \boxminus \rangle$ is an arbitrary but fixed basic revision-contraction pair. An interesting property of basic revision-contraction pairs is the following.

Proposition 1 *For all $K \in \mathcal{K}$ and $\phi \in L_*$, we have $\phi \notin K \boxminus \phi$ iff $\phi \notin K \boxplus \neg\phi$.*

Thus a sentence is retractible with respect to \boxminus from a given belief set K iff revising K , according to \boxplus , by its negation would lead us to reject it. Given this proposition and the basic postulate (r-Consistency) we can see that \boxminus satisfies (c-Success) whenever \boxplus satisfies (r-Success), although it can be shown the converse is false.⁸ The following proposition, in conjunction with Proposition 1 provides a counter-example to show that \boxminus does not generally satisfy (c-Success).

⁸Due to space limitations we omit the relevant counter-example in this version of the paper. Similar remarks apply at several other points in the paper.

Proposition 2 *There exists a belief negotiation model \mathcal{N} (such that $g_{\mathcal{N}}$ satisfies (g1) and (g2) – see Section 4) such that, for some $K \in \mathcal{K}$ and $\phi \in L_*$, we have $\phi \in K \boxplus_{\mathcal{N}} \neg\phi$.*

Note that, although (\boxplus, \boxminus) satisfies (1/2-Levi), the “other half” of the Levi Identity, i.e., $(K \boxminus \neg\phi) + \phi \subseteq K \boxplus \phi$, does *not* hold in general, since if it did, then \boxplus would satisfy (r-Success) (since $\phi \in (K \boxminus \neg\phi) + \phi$ for all ϕ). In fact it can be shown that for basic revision-contraction pairs this half of the Levi Identity is *equivalent* to (r-Success). As for the Harper Identity, well to begin with (Mixed Inclusion) (together with (r-Extensionality) and (c-Extensionality)) and (c-Inclusion) give us

$$K \boxminus \phi \subseteq (K \boxplus \neg\phi) \cap K \quad (1/2\text{-Harper } 1)$$

However, as in the case with the Levi Identity, we cannot in general strengthen this to equality. In fact the Harper Identity does not hold, in general, for *any* of the classes of basic revision-contraction pairs we present in this paper:

Proposition 3 *There exists a belief negotiation model \mathcal{N} (such that $g_{\mathcal{N}}$ satisfies (g3) – see Section 4) such that, for some $K \in \mathcal{K}$ and $\phi \in L_*$, we have $(K \boxplus_{\mathcal{N}} \neg\phi) \cap K \not\subseteq K \boxminus_{\mathcal{N}} \phi$.*

This counter-example together with the following proposition also shows that \boxminus does not in general satisfy (c-Recovery).

Proposition 4 \boxminus *satisfies (c-Recovery) iff the following two properties hold for all $K \in \mathcal{K}$ and $\phi \in L_*$:*

$$(K \boxplus \neg\phi) \cap K \subseteq K \boxminus \phi \quad (1/2\text{-Harper } 2)$$

$$K \subseteq (K \boxplus \neg\phi) + \phi \quad (r\text{-Recovery})$$

Note that (r-Recovery) follows from (r-Success). It can be shown that (r-Recovery) does not hold in general for basic revision-contraction pairs.

It is natural to ask what happens if we actually *force* the Harper Identity to hold. The following result shows the “harmlessness” of this exercise.

Proposition 5 *Let $\boxplus : \mathcal{K} \times L^* \rightarrow 2^L$ be a function which satisfies (r-Closure), (r-Extensionality), (r-Inclusion), (r-Vacuity) and (r-Consistency), and suppose $\boxminus : \mathcal{K} \times L_* \rightarrow 2^L$ is defined from \boxplus via the Harper Identity. Then (\boxplus, \boxminus) forms a basic revision-contraction pair.*

This result shows that there are in fact (at least) *two* ways we can define a basic revision-contraction pair from a given belief negotiation model \mathcal{N} . The first way is, as before, to define $(\boxplus_{\mathcal{N}}, \boxminus_{\mathcal{N}})$. The second way is to replace $\boxminus_{\mathcal{N}}$ here by $\boxminus_{\mathcal{N}}^H$, which is the contraction operator defined from $\boxplus_{\mathcal{N}}$ via the Harper Identity. By (1/2-Harper 1), we have $K \boxminus_{\mathcal{N}} \phi \subseteq K \boxminus_{\mathcal{N}}^H \phi$ for all $K \in \mathcal{K}$ and $\phi \in L_*$, and so, from the agent’s point of view, a case can perhaps be made for preferring the latter method, since it entails the agent giving up fewer beliefs when performing a contraction. Note that, by Proposition 1, $\boxminus_{\mathcal{N}}^H$ will satisfy (c-Success) iff $\phi \notin K \boxplus_{\mathcal{N}} \neg\phi$ for all K and $\phi \in L_*$ while, by Proposition 4, $\boxminus_{\mathcal{N}}^H$ will satisfy (c-Recovery) iff $\boxplus_{\mathcal{N}}$ satisfies (r-Recovery).

4 Restricting the belief negotiation model

Given a belief negotiation model $\mathcal{N} = \langle g, \{\nabla_\sigma\}_{\sigma \in \Sigma} \rangle$ there is clearly scope for making many restrictions on g and the ∇_σ — beyond the basic **(g0)** and **(v0)** — in order to capture different behaviour for $\boxplus_{\mathcal{N}}$ and $\boxminus_{\mathcal{N}}$, for example to try and capture weaker versions of the success postulates. In this paper we leave the ∇_σ alone and see what variations can be achieved just by fiddling with g . We select three restrictions in turn on g , starting with the mildest. We obtain three corresponding axiomatisations of the resulting classes of basic revision-contraction pairs.

4.1 First restriction

Consider first of all the following property of a function $g : \Sigma \rightarrow 2^{\mathcal{B}}$, for all $\sigma \in \Sigma$,

$$\text{(g1)} \quad |g(\sigma)| = 1$$

For a belief negotiation model \mathcal{N} , the function $g_{\mathcal{N}}$ will satisfy **(g1)** iff each negotiation round according to \mathcal{N} produces a winner, who does not make any concessions, and a loser who does. In other words forcing $g_{\mathcal{N}}$ to satisfy **(g1)** removes the possibility of a negotiation round ending in a draw. We may prove the following:

Theorem 2 *Let $\langle \boxplus, \boxminus \rangle$ be a basic revision-contraction pair. Then the following are equivalent:*

(i). \boxplus satisfies the following property, for all $K \in \mathcal{K}$ and $\phi \in L^*$,

If $\neg\phi \in K$ and $\neg\phi \notin K \boxplus \phi$ then $K \cup K \boxplus \phi$ is inconsistent

(Weak Consistent Expansion)

(ii). $\langle \boxplus, \boxminus \rangle = \langle \boxplus_{\mathcal{N}}, \boxminus_{\mathcal{N}} \rangle$ for some belief negotiation model \mathcal{N} such that $g_{\mathcal{N}}$ satisfies **(g1)**.

Note that the rule (Weak Consistent Expansion) is a weakening of (r-Success), which implies that $K \cup K \boxplus \phi$ is inconsistent *whenever* $\neg\phi \in K$. It can be shown this rule does not hold in general for basic revision-contraction pairs.

4.2 Second restriction

Now consider the following condition on a function $g : \Sigma \rightarrow 2^{\mathcal{B}}$:

(g2) Given $\sigma = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle) \in \Sigma$, if $S_i \in g(\sigma_i)$ for some $i < n$, then $S_n \in g(\sigma)$

This rule says that as soon as we reach a stage in which S_0 must be weakened, S_0 must be weakened further and further at *every* subsequent stage in the negotiation. In conjunction with **(g1)** this rule says that in a negotiation process between S and T *all* the weakening of T , if any, is done *before* any weakening of S takes place.

Theorem 3 Let $\langle \boxplus, \boxminus \rangle$ be a basic revision-contraction pair. Then the following are equivalent:

(i). $\langle \boxplus, \boxminus \rangle$ satisfies, for all $K \in \mathcal{K}$ and $\phi \in L^*$,

If $K \not\subseteq K \boxminus \neg\phi$ then $K \cup K \boxplus \phi$ is inconsistent (Consistent Retainment)

(ii). $\langle \boxplus, \boxminus \rangle = \langle \boxplus_{\mathcal{N}}, \boxminus_{\mathcal{N}} \rangle$ for some belief negotiation model \mathcal{N} such that $g_{\mathcal{N}}$ satisfies (g1) and (g2).

It can be shown that, for basic revision-contraction pairs, (Consistent Retainment) is a strictly stronger property than (Weak Consistent Expansion). The following proposition gives a way of re-expressing it.

Proposition 6 For basic revision-contraction pairs, the rule (Consistent Retainment) is equivalent to the conjunction of the following two rules:

If $K \subseteq K \boxplus \phi$ then $K \boxminus \neg\phi = K$ (Restricted Harper)

If $K \not\subseteq K \boxplus \phi$ then $K \cup K \boxplus \phi$ is inconsistent (Consistent Expansion)

The first rule above provides a restricted form of the Harper Identity. The second has already been studied in [3] as a characteristic postulate for selective revision operators. As is shown in that paper, it is a consequence of (r-Success) together with (r-Vacuity). The following result may be proved using Theorem 3 and Propositions 5 and 6. It provides an alternative characterisation of a particular class of selective revision operators which was characterised in [3] (Theorem 4.5 there, though note that the model of [3] also handles the limiting case of revising by an inconsistent sentence).

Theorem 4 Let $\boxplus : \mathcal{K} \times L^* \rightarrow 2^L$ be a function. Then the following are equivalent:

(i). \boxplus satisfies the rules (r-Closure), (r-Extensionality), (r-Inclusion), (r-Vacuity) (r-Consistency) and (Consistent Expansion).

(ii). $\boxplus = \boxplus_{\mathcal{N}}$ for some belief negotiation model \mathcal{N} such that $g_{\mathcal{N}}$ satisfies (g1) and (g2).

Note that, by Proposition 2, the two success postulates are still not satisfied. It can also be shown that (r-Recovery) does not generally hold for basic revision-contraction pairs satisfying (Consistent Retainment).

4.3 Last restriction

Finally we add our strongest restriction.

(g3) Given $\sigma = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle) \in \Sigma$, we have $g(\sigma) = \{S_n\}$

Note that (g3) implies both (g1) and (g2). This rule says that only S is weakened. T stays the same throughout. Thus any belief negotiation model \mathcal{N} such that $g_{\mathcal{N}}$ satisfies (g3) is heavily biased towards information source t . We may prove the following:

Theorem 5 *Let $\langle \boxplus, \boxminus \rangle$ be a basic revision-contraction pair. Then the following are equivalent:*

- (i). \boxplus satisfies (*r-Success*) (and so \boxminus satisfies (*c-Success*)).
- (ii). $\langle \boxplus, \boxminus \rangle = \langle \boxplus_{\mathcal{N}}, \boxminus_{\mathcal{N}} \rangle$ for some belief negotiation model \mathcal{N} such that $g_{\mathcal{N}}$ satisfies (**g3**).

Thus forcing $g_{\mathcal{N}}$ to satisfy (**g3**) leads $\boxplus_{\mathcal{N}}$ to satisfy *all* the basic AGM revision postulates. In other words $\boxplus_{\mathcal{N}}$ becomes a partial meet revision operator. The operator $\boxminus_{\mathcal{N}}$, meanwhile, now satisfies all the basic AGM contraction postulates with the possible exception (by Propositions 3 and 4) of (*c-Recovery*). Thus $\boxminus_{\mathcal{N}}$ becomes a withdrawal operator.

5 Related work on merging

As we said at the start of Section 2 the general problem of merging information coming from different sources has already received various treatments. In this section we briefly mention a few of these, some of which (namely [2, 11, 14]) actually generalise further than us in that they consider the problem of simultaneously merging together n pieces of information where possibly $n > 2$.

A difference between the present approach and papers such as [11, 12, 15] is that the latter are interested only in *fair* merging, i.e., they assume that the merging should always give equal precedence to the pieces of information involved, regardless of their source. In the case of binary merging which we consider this forces the merge operator to be commutative. Other approaches, such as [2, 14], drop this fairness requirement. Both these papers introduce explicit orderings of reliability or trustworthiness between sources (or, in the case of [2], between *sets* of sources), with precedence being given to information originating from the most reliable sources. In addition, [14] assumes that the information provided by a source takes the form of a total pre-order over the entire set \mathcal{W} which intuitively represents a grading of the worlds in \mathcal{W} according to their relative plausibility. This is a richer type of information which goes beyond just saying that the actual world is one of the worlds in a given subset of \mathcal{W} . The output of the merging operator of [14] is then another total pre-order over \mathcal{W} . (See also [10] for related work on "multi-agent belief revision"). It would be interesting to find out to what extent all this extra structure can be incorporated into our approach, maybe as a possible basis from which to define concrete instances of our functions g and ∇_{σ} . (Indeed the reader will notice that we already used the notion of source-reliability informally to explain Example 1 in Section 2.)

6 Conclusion and further work

We have defined a negotiation-style framework for merging two pieces of information coming from two distinct sources and have applied it to belief change. This framework supports the study of operations of contraction which do not

necessarily satisfy (c-Success) and (c-Recovery). We showed how a particular strain of Fermé and Hansson’s selective revision, as well as partial meet revision, can be captured by adding simple constraints on the function g of a belief negotiation model $\mathcal{N} = \langle g, \{\nabla_\sigma\}_{\sigma \in \Sigma} \rangle$ which basically relate to the order in which concessions are made. Other constraints of this type are possible. For example the following condition, in the presence of (g1), is a strengthening of (g2)

$$\text{Given } \sigma = (\langle S_0, T_0 \rangle, \dots, \langle S_n, T_n \rangle) \in \Sigma, S_n \in g(\sigma) \text{ iff } S_0 \in g(\sigma_0)$$

When taken together with (g1) this rule has the effect that the decision is made at the start of the negotiation which party must weaken, and then that party, and only that party, must continue weakening until the negotiation ends. Due to space limitations, the study of the effect of making this and other restrictions on g (not to mention possible restrictions on the ∇_σ) will have to wait. Another area of future research relates to the origin of the functions g and ∇_σ . In this paper we just assumed these functions were given. It would be interesting to define more concrete instances of them. Finally in this paper we represented the agent’s belief state as a belief set, i.e., a deductively closed set of sentences. Alternative ways of formally representing belief states have been proposed, most notably as *belief bases*, i.e., sets of sentences which are *not* necessarily deductively closed [8]. Operators of non-prioritised revision for belief bases have been studied [7, 9]. It should be possible to modify our framework so that it too handles this case. This too will be left to future work.

Acknowledgements

This work is supported by the DFG project “Computationale Dialektik”. Thanks are due to Gerhard Brewka, Daniela Carbogim, Sven Ove Hansson, Renata Wassermann and several anonymous referees for their helpful and encouraging comments.

References

- [1] C. Alchourrón, P. Gärdenfors and D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, *Journal of Symbolic Logic* 50 (1985) 510–530.
- [2] J. Cantwell, Resolving conflicting information, *Journal of Logic, Language and Information* 7(2) (1998) 191–220.
- [3] E. Fermé and S. O. Hansson, Selective revision, *Studia Logica* 63 (1999) 331–342.
- [4] E. Fermé and S. O. Hansson, Shielded contraction, in H. Rott and M. A. Williams (eds.), *Frontiers in Belief Revision*, Kluwer Academic Publishers, to appear.

- [5] A. K. Ghose and R. G. Goebel, Belief states as default theories: Studies in non-prioritized belief change, in *Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI'98)* (1998) 8–12.
- [6] S. M. Glaister, Recovery recovered, *Journal of Philosophical Logic*, **29** (2000) 171–206.
- [7] S. O. Hansson, Semi-Revision, *Journal of Applied Non-Classical Logic*, **7** (1997) 151–175.
- [8] S. O. Hansson, *A Textbook of Belief Dynamics*, Kluwer Academic Publishers (1999).
- [9] S. O. Hansson, A survey of non-prioritized belief revision, *Erkenntnis* **50** (1999) 413–427.
- [10] N. E. Kfir-Dahav and M. Tennenholtz, Multi-agent belief revision, in *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge (TARK'96)* (1996) 175–194.
- [11] S. Konieczny and R. Pino-Pérez, On the logic of merging, in *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)* (1998) 488–498.
- [12] P. Liberatore and M. Schaerf, Arbitration (or how to merge knowledge bases), *IEEE Transactions on Knowledge and Engineering* **10**(1) (1998) 76–90.
- [13] D. Makinson, On the status of the postulate of recovery in the logic of theory change, *Journal of Philosophical Logic*, **16** (1987) 383–394.
- [14] P. Maynard-Reid II and Y. Shoham, Belief fusion: Aggregating pedigreed belief states, *Journal of Logic, Language and Information*, **10**(2) (2001) 183–209.
- [15] P. Z. Revesz, On the semantics of theory change: Arbitration between old and new information, in *Proceedings of the Twelfth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS'93)* (1993) 71–82.
- [16] H. Rott, Preferential belief change using generalized epistemic entrenchment, *Journal of Logic, Language and Information* **1** (1992) 45–78.
- [17] D. N. Walton and E. C. W. Krabbe, *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, Albany, State University of New York Press (1995).
- [18] M. J. Wooldridge and S. Parsons, Languages for negotiation, in *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI2000)* (2000).