A Consistency-Based Framework for Merging Knowledge Bases

James P. Delgrande *

School of Computing Science, Simon Fraser University , Burnaby, B.C., Canada V5A 1S6

Torsten Schaub [1]

Institut für Informatik, Universität Potsdam, D–14415 Potsdam, Germany

---

Abstract

We present a framework for expressing various merging operators for belief sets. This framework generalises our earlier work on consistency-based belief revision and contraction. Two primary merging operators are identified: in the first approach, belief sources are consistently combined so that the result of merging knowledge bases $K_1, \ldots, K_n$ is a maximal consistent (if possible) set of formulas comprising the joint knowledge of the knowledge bases. This approach then accords with one's intuitions as to what a "merge" operator should do. The second approach is more akin to a generalised belief revision operator. Knowledge bases $K_1, \ldots, K_n$ are "projected" onto another (in the simplest case the knowledge base where only tautologies are known). Properties of these operators are investigated, primarily by comparing their properties with postulates that have been identified previously in the literature. Notably, the approach is independent of syntax, in that merging knowledge bases $K_1, \ldots, K_n$ is independent of how each K_i is expressed. As well, we investigate the role of entailment-based and consistency-based integrity constraints, the interrelationships between these approaches and belief revision, and the expression of further merging operators.

Key words: Belief set merging, consistency-based reasoning

---

* Corresponding author.
Email addresses: `jim@cs.sfu.ca` (James P. Delgrande), `torsten@cs.uni-potsdam.de` (Torsten Schaub).
URLs: `http://www.cs.sfu.ca/~jim` (James P. Delgrande), `http://www.cs.uni-potsdam.de/~torsten` (Torsten Schaub).

[1] Affiliated with the School of Computing Science at Simon Fraser University, Canada.

1 Introduction

The problem of merging multiple, potentially conflicting bodies of information arises in various guises. For example, an intelligent agent may receive reports from differing sources of knowledge that must be combined. As well, an agent may receive conflicting information from sensors that needs to be reconciled. Alternately, knowledge bases or databases comprising collections of data may need to be combined into a coherent whole. Even in dealing with a single, isolated, agent the problem of merging knowledge sets may arise: consider an agent whose beliefs are modelled by various independent "states of mind", but where it is desirable in some circumstances to combine such states of mind into a coherent whole, for example, before acting in a crucial situation. In all these cases, the fundamental problem is that of combining knowledge bases that may be mutually inconsistent, or conflicting, to get a coherent merged set of beliefs.

Given this diversity of situations in which the problem may arise, it is not surprising that different approaches have arisen for combining sources of information. The major subtypes of merging that have been proposed are called (following [1]) *majority* and *arbitration* operators. In the former case, the majority opinion counts towards resolving conflicts; in the latter, informally, the idea is to try to arrive at some consensus. In this paper, we develop a specific framework for defining merge operations. This framework extends our earlier work in belief revision. In both cases, the central intuition is that for belief change one begins by expressing the various knowledge bases, belief sources, etc. in distinct languages, and then (according to the belief change operation) in one way or another re-express the knowledge bases in a common language. Two approaches are first presented. In the first case, the intuition is that for merging knowledge bases, the common information is in a sense "pooled". This approach then seems to conform more naturally to the commonsense notion of merging of knowledge. A key property of this approach is that knowledge common to the knowledge bases is contained in the merged knowledge base. Thus if one knowledge base contained $p \wedge q$ and another $\neg p \wedge \neg q$, then $(p \wedge q) \vee (\neg p \wedge \neg q)$ would be in the merged knowledge base. Hence in this approach to merging, an intuition underlying the merging operation is that (at least) one of the knowledge bases contains correct information, but it is not known which.

In a second approach, knowledge bases are projected onto a separate knowledge base (which in the simplest case would consist solely of the set of tautologies). That is, the knowledge bases we wish to merge are used to augment the knowledge of a "target" body of knowledge. This second approach then appears to be a natural extension of belief *revision*. In this approach, knowledge common to the knowledge bases may not be contained in the merged knowledge base. Thus if two knowledge bases contained $p \wedge q$ and $\neg p \wedge \neg q$, respectively, then $(p \wedge q) \vee (\neg p \wedge \neg q)$ may not be in the merged knowledge base; for example $p \wedge \neg q$ may be consistent with the merged knowledge base. Hence, here an intuition underlying the merging operation is that perhaps some "common ground" is found between the merged knowledge bases.

In both approaches, we address the role of entailment-based and consistency-based integrity constraints with respect to the merge operators. Both approaches have reasonable properties, compared with postulate sets that have appeared in the literature. As well, the second type of approach has not, to our knowledge, been investigated previously. The next section describes related work while Section 3 develops these approaches. Following this, we consider variants on these approaches, including prioritised merging. Section 5 briefly considers computational complexity. We conclude with a discussion. Proofs of theorems are found in an appendix.

2 Background

2.1 Consistency-Based Belief Revision

This subsection summarises our earlier work [2] on consistency-based belief revision. Throughout this paper, we deal with propositional languages and use the logical symbols $\top$, $\bot$, $\neg$, $\vee$, $\wedge$, $\supset$, and $\equiv$ to construct formulas in the standard way. We write $\mathcal{L}_{\mathcal{P}}$ to denote a language over an alphabet $\mathcal{P}$ of *propositional letters* or *atomic propositions*. Formulas are denoted by the Greek letters α, β, α_1, ... *Knowledge bases* are identified with deductively-closed sets of formulas, or *belief sets*, and are denoted K, K_1, ...[2] Thus $K = Cn(K)$, where $Cn(\cdot)$ is the deductive closure in classical propositional logic of the formula or set of formulas given as argument. Given an alphabet $\mathcal{P}$, we define a disjoint alphabet $\mathcal{P}'$ as $\mathcal{P}' = \{p' \mid p \in \mathcal{P}\}$. For $\alpha \in \mathcal{L}_{\mathcal{P}}$, α' is the result of replacing in α each proposition $p \in \mathcal{P}$ by the corresponding proposition $p' \in \mathcal{P}'$ (so implicitly there is an isomorphism between $\mathcal{P}$ and $\mathcal{P}'$, and thus $\mathcal{L}_{\mathcal{P}}$ and $\mathcal{L}_{\mathcal{P}'}$). This is defined analogously for sets of formulas. This notation essentially allows us to refer to a formula or set of formulas relative to a knowledge base. In turn, this means that we can rely on the fact that, while p and $\neg p$ are mutually contradictory, p and $\neg p'$, trivially, are not.

A *belief change scenario* in $\mathcal{L}_{\mathcal{P}}$ is a triple $B = (K, R, C)$ where K, R, and C are sets of formulas in $\mathcal{L}_{\mathcal{P}}$. Informally, K is a belief set that is to be modified so that the formulas in R are contained in the result, and the formulas in C are not. For an approach to revision we have $|R| = 1$ and $C = \emptyset$, and for an approach to contraction we have $R = \emptyset$ and $|C| = 1$. An extension determined by a belief change scenario, called a *belief change extension*, is defined as follows.

[2] We note that while we deal solely with belief sets in this paper, our definitions work for arbitrary sets of formulas.

Definition 2.1 *Let $B = (K, R, C)$ be a belief change scenario in $\mathcal{L}_\mathcal{P}$.*

Define EQ as a maximal set of equivalences $EQ \subseteq \{p \equiv p' \mid p \in \mathcal{P}\}$ such that

$$Cn(K' \cup R \cup EQ) \cap (C \cup \{\perp\}) \;=\; \emptyset.$$

Then $Cn(K' \cup R \cup EQ) \cap \mathcal{L}_\mathcal{P}$ *is a* (consistent) belief change extension *of B.*

If there is no such set EQ then B is inconsistent *and $\mathcal{L}_\mathcal{P}$ is defined to be the sole* (inconsistent) belief change extension *of B.*

Note that in the definition, "maximal" is with respect to set containment (rather than set cardinality). The exclusive use of "$\{\perp\}$" in the definition is to take care of consistency if $C = \emptyset$. Clearly a consistent belief change extension of B is a modification of K which contains every formula in R, and which contains no formula in C. We say that EQ *determines* the respective consistent belief change extension of B. For a given belief change scenario there may be more than one consistent belief change extension. We make use of the notion of a *selection function* c that for any set $I \neq \emptyset$ has as value some element of I. In defining revision, we use such a selection function to select a specific consistent belief change extension. [3]

Definition 2.1 provides a very general framework for specifying belief change. We can restrict the definition to obtain specific functions for belief revision and contraction; here we just deal with revision.

Definition 2.2 (Revision) *Let K be a belief set and α a formula, and let $(E_i)_{i\in I}$ be the family of all belief change extensions of $(K, \{\alpha\}, \emptyset)$. Then, we define*

1. $K \dot{+}_c \alpha = E_i$ *as a* choice revision *of K by α with respect to some selection function c with $c(I) = i$.*
2. $K \dot{+} \alpha = \bigcap_{i\in I} E_i$ *as the* (skeptical) revision *of K by α.*

For instance, (skeptically) revising $Cn(p \wedge q)$ by $\neg q$ results in $Cn(p \wedge \neg q)$. This belief change extension is determined by $\{p \equiv p'\}$ from the renamed belief set $\{p' \wedge q'\}$ and the revision formula $\neg q$. As a second example, we get

$$\{\neg p \equiv q\} \dot{+} \neg q \;=\; Cn(p \wedge \neg q)$$

from the renamed knowledge base $\neg p' \equiv q'$ and formula $\neg q$, along with equivalences $\{p \equiv p', q \equiv q'\}$. For a third example, observe that both $\{p \vee q\} \dot{+} (\neg p \vee \neg q)$ as well as $\{p \wedge q\} \dot{+} (\neg p \vee \neg q)$ result in $Cn(p \equiv \neg q)$, although the former is de-

[3] This use of selection functions is slightly different from that in the AGM approach.

termined by $\{p \equiv p', q \equiv q'\}$, while the latter relies on two such sets, viz. $\{p \equiv p'\}$ and $\{q \equiv q'\}$.

With respect to related work, and specifically to the foundational AGM postulates [3], we obtain that the basic postulates are satisfied, along with supplementary postulate $(K\dot{+}7)$ for both choice and skeptical revision.

Definition 2.1 also leads to a natural and general treatment of both consistency-based and entailment-based integrity constraints; see [2] for details.

2.2 *Belief Merging*

In this section we review related work in belief set merging. We focus on two sets of postulates that have been used to characterise merging, and with respect to which we compare our own approaches. Following this we briefly survey representative related work in the literature.

First, Liberatore and Schaerf [4] consider merging two belief bases and propose the following postulate set to characterise a merge operator that they call an *arbitration* operator ([5] call this a *commutative revision operator*). They restrict their attention to propositional languages over a finite set of atoms; consequently their merging operator can be expressed as a binary operator on formulas. They provide the following postulates, which we express as a definition.

Definition 2.3 *$\triangle$ is an arbitration operator (or commutative revision operator) iff $\triangle$ satisfies the following postulates.*

$(LS1)$ $\vdash \alpha \triangle \beta \equiv \beta \triangle \alpha$.
$(LS2)$ $\vdash \alpha \wedge \beta \supset \alpha \triangle \beta$.
$(LS3)$ *If $\alpha \wedge \beta$ is satisfiable then $\vdash \alpha \triangle \beta \supset \alpha \wedge \beta$.*
$(LS4)$ *$\alpha \triangle \beta$ is unsatisfiable iff α is unsatisfiable and β is unsatisfiable.*
$(LS5)$ *If $\vdash \alpha_1 \equiv \alpha_2$ and $\vdash \beta_1 \equiv \beta_2$ then $\vdash \alpha_1 \triangle \beta_1 \equiv \alpha_2 \triangle \beta_2$.*
$(LS6)$ $\alpha \triangle (\beta_1 \vee \beta_2) = \begin{cases} \alpha \triangle \beta_1 & \textit{or} \\ \alpha \triangle \beta_2 & \textit{or} \\ (\alpha \triangle \beta_1) \vee (\alpha \triangle \beta_2) & \end{cases}$
$(LS7)$ $\vdash (\alpha \triangle \beta) \supset (\alpha \vee \beta)$.
$(LS8)$ *If α is satisfiable then $\alpha \wedge (\alpha \triangle \beta)$ is satisfiable.*

The first postulate asserts that the merging is a commutative operator, while the next two assert that, for mutually consistent formulas, merging corresponds to their conjunction. $(LS5)$ ensures that the operator is independent of syntax, while $(LS6)$ provides a “factoring” postulate, analogous to a similar factoring result in (AGMstyle) belief revision and contraction. Postulate $(LS7)$ can be taken as distinguishing $\triangle$ from other such operators; it asserts that the result of merging implies the disjunction of the original formulas. The last postulate informally constrains the result of merging so that each operator "contributes to" (i.e. is consistent with) the final result.

Konieczny and Pino Peréz [5] also consider the problem of merging possibly contradictory belief bases. To this end, they consider finite multisets of the form $\Psi = \{K_1, \ldots, K_n\}$ and assume that all belief sets K_i are consistent, finitely representable, and therefore representable by a formula. K^{+n} is the multiset consisting of n copies of K.[4] Multiset union is denoted $\sqcup$, wherein for example $\{\phi\} \sqcup \{\phi\} = \{\phi, \phi\}$. Following [5], we use[5] $\Delta^{\mu}(\Psi)$ to denote the result of merging the multiset Ψ of belief bases given the entailment-based integrity constraint expressed by μ. They provide the following set of postulates:

Definition 2.4 ([5]) *Let Ψ be a multiset of sets of formulas, and ϕ, μ formulas (all possibly subscripted or primed). Δ is an* IC merging operator *iff it satisfies the following postulates.*

$(IC0)$ $\Delta^{\mu}(\Psi) \vdash \mu$.
$(IC1)$ *If* $\mu \nvdash \bot$ *then* $\Delta^{\mu}(\Psi) \nvdash \bot$.
$(IC2)$ *If* $\bigwedge \Psi \nvdash \neg\mu$ *then* $\Delta^{\mu}(\Psi) \equiv \bigwedge \Psi \wedge \mu$.
$(IC3)$ *If* $\Psi_1 \equiv \Psi_2$ *and* $\mu_1 \equiv \mu_2$ *then* $\Delta^{\mu_1}(\Psi_1) \equiv \Delta^{\mu_2}(\Psi_2)$.
$(IC4)$ *If* $\phi \vdash \mu$ *and* $\phi' \vdash \mu$ *then:* $\Delta^{\mu}(\phi \sqcup \phi') \wedge \phi \nvdash \bot$ *implies* $\Delta^{\mu}(\phi \sqcup \phi') \wedge \phi' \nvdash \bot$.
$(IC5)$ $\Delta^{\mu}(\Psi_1) \wedge \Delta^{\mu}(\Psi_2) \vdash \Delta^{\mu}(\Psi_1 \sqcup \Psi_2)$.
$(IC6)$ *If* $\Delta^{\mu}(\Psi_1) \wedge \Delta^{\mu}(\Psi_2) \nvdash \bot$ *then* $\Delta^{\mu}(\Psi_1 \sqcup \Psi_2) \vdash \Delta^{\mu}(\Psi_1) \wedge \Delta^{\mu}(\Psi_2)$.
$(IC7)$ $\Delta^{\mu_1}(\Psi) \wedge \mu_2 \vdash \Delta^{\mu_1 \wedge \mu_2}(\Psi)$.
$(IC8)$ *If* $\Delta^{\mu_1}(\Psi) \wedge \mu_2 \nvdash \bot$ *then* $\Delta^{\mu_1 \wedge \mu_2}(\Psi) \vdash \Delta^{\mu_1}(\Psi) \wedge \mu_2$.

The intent is that $\Delta^{\mu}(\Psi)$ is the belief base closest to the belief multiset Ψ. Of the postulates, $(IC2)$ states that the result of merging is simply the conjunction of the belief bases and integrity constraints, when consistent. $(IC4)$ is a *fairness* postulate, that when two belief bases disagree, merging doesn't give preference to one of them. $(IC5)$ states that a model of two mergings is in the union of their merging. With $(IC5)$ we get that if two mergings are consistent then their merging is implied by their conjunction. Note that merging operators are trivially commutative. $(IC7)$ and $(IC8)$ correspond to the extended AGM postulates $(K\dot{+}7)$ and $(K\dot{+}8)$ for revision, but with respect to the integrity constraints. Postulates $(IC1)$–$(IC6)$, with tautologous integrity constraints, correspond to basic merging without integrity constraints in [1].

---

[4] [5] uses the notation K^n.
[5] [5] writes $\Delta_{\mu}(\Psi)$ where we have $\Delta^{\mu}(\Psi)$.

A majority operator is characterised in addition by the postulate:

$(Maj)\ \exists n \Delta^{\mu}(\Psi_1 \sqcup \Psi_2^{+n}) \vdash \Delta^{\mu}(\Psi_2)$

Thus, given enough repetitions of a belief base Ψ_2, this belief base will eventually come to dominate the merge operation.

An arbitration operator is characterised by the original postulates together with the following postulate; see [5] for an explanation.

(Arb) Let μ_1 and μ_2 be logically independent. If $\Delta^{\mu_1}(\phi_1) \equiv \Delta^{\mu_2}(\phi_2)$ and $\Delta^{\mu_1 \equiv \mu_2}(\phi_1 \sqcup \phi_2) \equiv (\mu_1 \equiv \mu_2)$ then $\Delta^{\mu_1 \vee \mu_2}(\phi_1 \sqcup \phi_2) \equiv \Delta^{\mu_1}(\phi_1)$.

[1] characterises these approaches as trying to minimize global dissatisfaction vs. trying to minimize local dissatisfaction respectively. Examples are given of a merging operator using Dalal's notion of distance [6].

Earlier work on merging operators includes [7] and [8]. The former proposes various theory merging operators based on the selection of maximum consistent subsets in the union of the belief bases; see [9] for a pertinent discussion. The latter proposes an "arbitration" operator that satisfies a subset of the Liberatore and Schaerf postulates; see [10] for a discussion. [11] first identified and addressed the majority merge operator. [12] gives a framework for defining merging operators, where a family of merging operators is parameterised by a distance between interpretations and aggregating functions. The authors suggest that most, if not all, model-based merging operators can be captured in their approach, along with a selection of syntax-based operators.

More or less concurrently, [13] proposed a general approach to formulating merging functions, based on ordinal conditional functions [14]. Roughly, epistemic states are associated with a mapping from possible worlds onto the set of ordinal numbers. Various merging operators then can be defined by considering the ways in which the "Cartesian product" of two epistemic states can be resolved into an ordinal conditional function. [15] also considers the problem of an agent merging information from different sources, via what is called *social contraction*. In a manner analogous to the Levi Identity for belief revision, information from the various sources is weakened to the extent that it can be consistently added to the agent's belief base. Last, much work has been carried out in merging possibilistic knowledge bases; see for example [16].

3 Consistency-Based Approaches to Belief Set Merging

In this section we modify the framework given by Definition 2.1 to deal with belief set *merging*, in which multiple sources of information (knowledge bases, etc.) are coalesced into a single belief set. We detail two different approaches to belief set merging in this section, expressible in the general approach.

In the first case, the intuition is that for merging belief sets, the common information is in a sense “pooled”. This approach then seems to conform to the commonsense notion of merging of knowledge, in which sets of knowledge are joined to produce a single knowledge set retaining as much as possible of the contents of the original knowledge sets. As well, it adheres for the most part to the Liberatore and Schaerf postulates [4]. In the second approach, knowledge sources are projected onto a separate knowledge source (which in the simplest case could consist solely of tautologies). That is, the sources we wish to merge are used to augment the knowledge of another source, which could be thought of as a set of integrity constraints or alternatively as a set of formulas for revision. This approach generally follows the postulates given by Konieczny and Pino Peréz [5].

3.1 Multi belief change scenarios

Our approaches to merging are centred around the notion of a *multi belief change scenario*:

Definition 3.1 *A* multi belief change scenario*, B, in $\mathcal{L}_{\mathcal{P}}$ is a triple*

$$B = (\mathcal{K}, R, C)$$

where $\mathcal{K}$ is a family $(K_j)_{j\in J}$ of sets of formulas in $\mathcal{L}_{\mathcal{P}}$, and $R, C \subseteq \mathcal{L}_{\mathcal{P}}$.

Informally, $\mathcal{K}$ is a collection of belief sets that are to be merged so that the formulas in R are contained in the result, and the formulas in C are not. So this is the same as a belief change scenario as defined in Section 2, except that the single set of formulas K is extended to several of sets of formulas. R and C will be used to express entailment-based and consistency-based integrity constraints, respectively. That is, the formulas in R will all be true in the result of a merging, whereas the formulas in C will not be contained in the result. While R is intended to represent a set of entailment-based integrity constraints [17], it could just as easily be regarded as a set of formulas for revision. Similarly, while C is intended to represent a set of (negations of) consistency-based integrity constraints [18,19], it could just as easily be regarded as a set of formulas for contraction. Thus the overall approaches can be considered as a framework in which merging, revising, and (multiple) contractions may be carried out in parallel while taking into account integrity constraints.

To begin with, we generalise the notation α' from Section 2 in the obvious way for integers $i > 0$ and sets of integers: for alphabet $\mathcal{P}$, we define $\mathcal{P}^i$ as $\mathcal{P}^i = \{p^i \mid p \in \mathcal{P}\}$, and α^i etc. analogous to Section 2. Similarly we define for a set or list of positive integers N that $\mathcal{P}^N = \{p^i \mid p \in \mathcal{P}, i \in N\}$. Then $\alpha^N = \{\alpha^i \mid i \in N\}$. The definition of an extension to a multi belief change scenario will depend on the specific approach to merging that is being formalised. We consider each approach in turn in the following two subsections.

3.2 Belief Set Merging

Consider the first approach, in which the contents of belief sets are to be merged. The issue of integrity constraints is addressed after the basic definitions are given.

Definition 3.2 *Let $B = (\mathcal{K}, \emptyset, \emptyset)$ be a multi belief change scenario in $\mathcal{L}_{\mathcal{P}}$, where $\mathcal{K} = (K_j)_{j \in J}$. Define EQ as a maximal set of equivalences*

$$EQ \subseteq \{p^k \equiv p^l \mid p \in \mathcal{P} \text{ and } k, l \in J\}$$

such that

$$Cn\Big(\textstyle\bigcup_{j \in J} K_j^j \cup EQ\Big) \cap \{\bot\} = \emptyset$$

Then

$$\Big\{\alpha \;\Big|\; \{\alpha^j \mid j \in J\} \subseteq Cn\Big(\textstyle\bigcup_{j \in J} K_j^j \cup EQ\Big)\Big\}$$

is a consistent symmetric belief change extension *of B.*

If there is no such set EQ then B is inconsistent *and $\mathcal{L}_{\mathcal{P}}$ is defined to be the sole* (inconsistent) symmetric belief change extension *of B.*

Definition 3.3 (Merging) *Let $\mathcal{K}$ be a family of sets of formulas in $\mathcal{L}_{\mathcal{P}}$ and let $(E_i)_{i \in I}$ be the family of all symmetric belief change extensions of $(\mathcal{K}, \emptyset, \emptyset)$.*

Then, we define

1. $\Delta_c(\mathcal{K}) = E_i$ *as the* choice merging *of $\mathcal{K}$ with respect to selection function c with $c(I) = i$.*
2. $\Delta(\mathcal{K}) = \bigcap_{i \in I} E_i$ *as the* (skeptical) merging *of $\mathcal{K}$.*

Of particular interest is *binary* merging, where $\mathcal{K} = \{K_1, K_2\}$. In this case, we will write the merge operator Δ as an infix operator. That is, $\Delta(\{K_1, K_2\})$ is written as $K_1 \vartriangle K_2$. Also, given two formulas α, β, we just write $\alpha \vartriangle \beta$.

Example 1 *$(p \wedge q \wedge r) \vartriangle (p \wedge \neg q \wedge s)$ yields (informally) $(p^1 \wedge q^1 \wedge r^1) \wedge (p^2 \wedge \neg q^2 \wedge s^2)$ along with $EQ = \{p^1 \equiv p^2, r^1 \equiv r^2, s^1 \equiv s^2\}$. The result of merging is $Cn(\{p \wedge r \wedge s\})$.*

Example 2 *Let*

$$K_1 \equiv p \wedge q \wedge r \wedge s \;\; \text{and} \;\; K_2 \equiv \neg p \wedge \neg q \wedge \neg r \wedge \neg s.$$

We obtain that $K_1 \vartriangle K_2$ yields $EQ = \emptyset$ and in fact

$$K_1 \vartriangle K_2 \; = \; Cn(\{(p \wedge q \wedge r \wedge s) \vee (\neg p \wedge \neg q \wedge \neg r \wedge \neg s)\}) \, .$$

This example is introduced and discussed in [1]; as well it corresponds to the postulate $(LS7)$. Consider where K_1 and K_2 represent two analyst's forecasts concerning how four different stocks are going to perform. p represents the fact that the first stock will rise, etc. The result of merging is a belief set, in which it is believed that either all will rise, or that all will not rise. That is, essentially, it is believed that one forecast will hold in its entirety, or the other will. As [1] points out, knowing nothing else and assuming independence of the stock's movements, this is implausible: it is possible that some stocks rise while others do not. On the other hand, if we have reason to believe that one analyst is in fact highly reliable (although we don't know which) then the result of Example 2 *is* reasonable. However this example illustrates that there are cases wherein this formulation is too strong.

We obtain the following with respect to the postulate sets described in Section 2.2. [6]

Theorem 3.1 *Let Δ and Δ_c be defined as in Definition 3.3, but letting μ, μ_1, and μ_2 be $\top$.* [7]

Then Δ and Δ_c satisfy the postulates $(IC0)$, $(IC2)$, $(IC3)$, $(IC5)$, $(IC7)$, $(IC8)$, as well as the weaker versions of $(IC1)$ [8] *and $(IC6)$, and a stronger version of $(IC4)$:*

$(IC1')$ *If $K \nvdash \bot$ for every $K \in \mathcal{K}$, then $\Delta(\mathcal{K}) \nvdash \bot$.*
$(IC4')$ *If $K \nvdash \bot$ for every $K \in \mathcal{K}$, then for $K \in \mathcal{K}$ we have $\Delta(\mathcal{K}) \cup K \nvdash \bot$.*
$(IC6')$ *If $K_1 \wedge K_2 \nvdash \bot$ then $\Delta(\{K_1\} \sqcup \{K_2\}) \vdash K_1 \wedge K_2$.*

[6] In discussing the IC postulates we will use the notation of [1]; for the LS postulates we will use the notation of [4].
[7] Setting μ, μ_1, and μ_2 to $\top$ reflects the fact that in Definition 3.1 we have $R = C = \emptyset$. For more explanation, see the discussion on integrity constraints following Theorem 3.3.
[8] It is straightforward to obtain $(IC1)$ by essentially ignoring inconsistent belief sets. We remain with the present postulate since it reflects the most natural formulation of merging in our framework.

A counterexample to $(IC6)$ is given by

$$\mathcal{K}_1 = \{Cn(p)\,,\ Cn(\neg p)\}, \qquad \mathcal{K}_2 = \{Cn(p)\}.$$

We have $\Delta(\mathcal{K}_1) \wedge \Delta(\mathcal{K}_2) \equiv \top \wedge p \equiv p$, while $\Delta(\mathcal{K}_1 \sqcup \mathcal{K}_2) \equiv \top$. The fact that this approach fails to satisfy $(IC6)$ as originally given seems reasonable to us, at least in the context of a non-majority merging operator. Indeed, it proves to be the case that present approach satisfies a non-majority postulate, viz.:

$$\Delta(\mathcal{K}_1 \sqcup \mathcal{K}_2^{+n}) \ = \ \Delta(\mathcal{K}_1 \sqcup \mathcal{K}_2).$$

This postulate is identified in [1], a weaker version of which is used to define their arbitration operator.

Theorem 3.2 *Let* Δ *and* Δ_c *be defined as in Definition 3.3.*

Then Δ *and* Δ_c *satisfy the following postulates.*

(1) $(LS1)$, $(LS2)$, $(LS3)$, $(LS5)$, $(LS7)$

as well as the following weaker versions of the remaining postulates:

(2) $(LS4)'$ $\alpha \vartriangle \beta$ *is satisfiable iff* α *is satisfiable and* β *is satisfiable.*
$(LS6)'$ $(\alpha \vartriangle \beta_1) \wedge \beta_2$ *implies* $\alpha \vartriangle (\beta_1 \wedge \beta_2)$.
$(LS8)'$ *If* α *is satisfiable and* β *is satisfiable then* $\alpha \wedge (\alpha \vartriangle \beta)$ *is satisfiable.*

(3) $(LS6c)'$ *For any selection function* c *there is a selection function* c' *such that* $\alpha \vartriangle_c \beta_1$ *implies* $\alpha \vartriangle_{c'} (\beta_1 \vee \beta_2)$ *or* $\alpha \vartriangle_c \beta_2$ *implies* $\alpha \vartriangle_{c'} (\beta_1 \vee \beta_2)$.

Example 3 *A counterexample to* $(LS6)$ *is given by the following.*

$$\alpha \ \equiv \ (p \wedge q \wedge r \wedge s)\,, \qquad \beta_1 \ \equiv \ (\neg p \wedge \neg q) \vee \neg r\,, \qquad \beta_2 \ \equiv \ \neg q \vee \neg s\,.$$

We get that:

$$\begin{aligned} \alpha \vartriangle (\beta_1 \vee \beta_2) \ &\equiv \ (p \wedge q \wedge r) \vee (p \wedge q \wedge s) \vee (p \wedge r \wedge s)\,, \\ \alpha \vartriangle \beta_1 \ &\equiv \ (p \wedge q \wedge s) \vee (r \wedge s)\,, \\ \alpha \vartriangle \beta_2 \ &\equiv \ (p \wedge q \wedge r) \vee (p \wedge r \wedge s)\,. \end{aligned}$$

This result echoes similar results in distance-based belief revision. Typically, AGM revision postulate $(K\dot{+}8)$ fails in such distance-based approaches [20]. Here, $(LS6)'$ corresponds to AGM revision postulate $(K\dot{+}7)$, while $(LS6)$ is analogous to a "factoring result" in revision that in turn is equivalent to $(K\dot{+}7)$ and $(K\dot{+}8)$.

While the merging operator is commutative by definition, it is not associative; for example $(((p \vee q) \vartriangle \neg p) \vartriangle p) \neq (p \vee q) \vartriangle (\neg p \vartriangle p)$. Lastly, we have the following result showing that in this approach, merging two belief sets is expressible in terms of our approach to revision.

Theorem 3.3 *Let $\dot{+}$ and $\vartriangle$ be given as in Definitions 2.2 and 3.3 (respectively). Then,*

$$\alpha \vartriangle \beta \;=\; \alpha \dot{+} \beta \cap \beta \dot{+} \alpha.$$

As in [2], we can also consider the role of integrity constraints in belief set merging. However there is a fundamental problem here: given the presence of postulate $(LS7)$, for (entailment-based) integrity constraint μ and merge operation $\alpha \vartriangle \beta$, it is unclear what the result should be when $\alpha \vee \beta \vdash \neg\mu$. The simplest solution is to simply have the result of merging be the inconsistent belief set when this occurs. Hence in [21] we had the following definition, in place of Definition 3.2. (Other definitions are unchanged.)

Definition 3.4 *Let $B = (\mathcal{K}, R, C)$ be a multi belief change scenario in $\mathcal{L}_{\mathcal{P}}$, where $\mathcal{K} = (K_j)_{j \in J}$. Define EQ as a maximal set of equivalences*

$$EQ \subseteq \{p^k \equiv p^l \mid p \in \mathcal{P} \text{ and } k, l \in J\}\}$$

such that

$$Cn\Big(\bigcup_{j \in J} K_j^j \cup R^J \cup EQ\Big) \cap (C^J \cup \{\bot\}) = \emptyset$$

Then

$$\Big\{\alpha \;\Big|\; \{\alpha^j \mid j \in J\} \subseteq Cn\Big(\bigcup_{j \in J} K_j^j \cup R^J \cup EQ\Big)\Big\}$$

is a consistent symmetric belief change extension *of B with integrity constraints.*

If there is no such set EQ then B is inconsistent *and $\mathcal{L}_{\mathcal{P}}$ is defined to be the sole* (inconsistent) symmetric belief change extension *of B.*

The sets R^J ensure that the integrity constraints in R are true in each belief set, and so will be true in the result. Of course, this may come at the cost of inconsistency for the merged belief set.

Otherwise there are two ways that one can ensure that the result of merging is consistent with an (entailment-based) integrity constraint: For multiset $\mathcal{K}$ and integrity constraint μ, one can ensure that each belief set in $\mathcal{K}$ is consistent with μ by revising by μ. Alternatively, one can merge the members of $\mathcal{K}$ and then revise by μ. Thus in the first case we define the merging of $\mathcal{K} = \{K_1, \ldots, K_n\}$ with constraint μ as

$$\Delta((K_j \dot{+} \mu)_{j \in J}).$$

In the second case we define the merging as

$$(\Delta(K_j)_{j \in J}) \dot{+} \mu.$$

These approaches are not equivalent, as the next example illustrates.

Example 4 *Let $K_1 = Cn(p)$, $K_2 = Cn(p \supset q)$, and $\mu = \neg p \vee \neg q$.*

Then $K_1 \vartriangle K_2 = Cn(p \wedge q)$ and $(K_1 \vartriangle K_2) \dot{+} \mu = Cn(p \equiv \neg q)$.

However, $K_1 \dot{+} \mu = Cn(p \wedge \neg q)$, $K_2 \dot{+} \mu = Cn(\neg p)$, and

$$(K_1 \dot{+} \mu) \vartriangle (K_2 \dot{+} \mu) = Cn(\neg q).$$

This example also shows that neither possibility is strictly stronger than the other.

3.3 Belief Set Projection

In our second approach, the contents of several belief sets are "projected" onto another designated belief set. Again, the formulation is straightforward within the framework of belief change scenarios. For belief sets $K_1, \ldots, K_n$, we express each in a distinct language, but project these belief sets onto a distinguished belief set in which R is believed. (In the simplest case we would have $R \equiv \top$.)

In the following, R and C represent a set of entailment-based and consistency-based integrity constraints, respectively.

Definition 3.5 *Let $B = (\mathcal{K}, R, C)$ be a multi belief change scenario in $\mathcal{L}_{\mathcal{P}}$, where $\mathcal{K} = (K_j)_{j \in J}$. Define EQ as a maximal set of equivalences*

$$EQ \subseteq \{p^j \equiv p \mid p \in \mathcal{P} \text{ and } j \in J\}$$

such that

$$Cn\Big(\textstyle\bigcup_{j \in J} K_j^j \cup R \cup EQ\Big) \cap (C \cup \{\bot\}) = \emptyset.$$

Then

$$Cn\Big(\bigcup_{j\in J} K_j^j \cup R \cup EQ\Big) \cap \mathcal{L}_{\mathcal{P}}$$

is a consistent projected belief change extension *of* B.

If there is no such set EQ *then* B *is* inconsistent *and* $\mathcal{L}_{\mathcal{P}}$ *is defined to be the sole* (inconsistent) projected belief change extension *of* B.

There is an interesting similarity between revision and projection. Revision in some sense "projects" a belief set onto the formula that we revise with. Similarly, the actual projection operation "projects" a family of belief sets onto whatever is contained in R.

Definition 3.6 (Merging via Projection) *Let* $\mathcal{K}$ *be a family of sets of formulas in* $\mathcal{L}_{\mathcal{P}}$*, let* R *and* C *be sets of formulas in* $\mathcal{L}_{\mathcal{P}}$*, and let* $(E_i)_{i\in I}$ *be the family of all projected belief change extensions of* $(\mathcal{K}, R, C)$.

Then, we define

1. $\nabla_c^{R,C}(\mathcal{K}) = E_i$ *as the* choice merging *of* $\mathcal{K}$ *with respect to integrity constraints* R *and* C*, and selection function* c *with* $c(I) = i$.
2. $\nabla^{R,C}(\mathcal{K}) = \bigcap_{i\in I} E_i$ *as the* (skeptical) merging *of* $\mathcal{K}$ *with respect to integrity constraints* R *and* C.

As above, for two formulas α and β, we just write $\alpha\nabla\beta$, if $R = C = \emptyset$ and we write $\alpha\nabla^{\mu}\beta$ if $R = \{\mu\}$ and $C = \emptyset$.

Example 5 *We have that* $(p \wedge q \wedge r)\nabla(p \wedge \neg q)$ *yields two* EQ *sets:*

$$EQ_1 = \{p^1 \equiv p,\ p^2 \equiv p,\ q^1 \equiv q,\ r^1 \equiv r,\ r^2 \equiv r\} \quad \textit{and}$$
$$EQ_2 = \{p^1 \equiv p,\ p^2 \equiv p,\ q^2 \equiv q,\ r^1 \equiv r,\ r^2 \equiv r\}.$$

The result of merging is $p \wedge r \wedge s$.

Example 6 *Consider the example from [1]:*

$$K_1 \equiv p \wedge q \wedge r \wedge s \ \textit{ and } \ K_2 \equiv \neg p \wedge \neg q \wedge \neg r \wedge \neg s.$$

In forming a set of equivalences, EQ*, we can have precisely one of* $p^1 \equiv p$ *or* $p^2 \equiv p$ *in* EQ*, and similarly for the other atomic sentences. Each such set of equivalences then represents one way each forecaster's prediction for a specific stock* *can be taken into account. Taken all together then we have 2^4 sets of equivalences, and in the end we obtain that*

$$K_1 \nabla K_2 \ = \ Cn(\top)\,.$$

We feel that this is a plausible outcome in the interpretation involving the forecasted movement of independent stocks. Note that if the example were extended so that multiple possibilities for stock movement were allowed, then we would obtain in the projection the various compromise positions for the two belief sets. Thus for example if a stock could either remain the same, or go up or down a little or a lot, and one forecaster predicted that stocks a and b would go up a lot, and another predicted that they would both go down a lot, then the projection would have both stocks moving a lot, although it would be unclear as to whether the movement would be up or down.

We obtain the following.

Theorem 3.4 *Let ∇ and ∇_c be defined as in Definition 3.6.*

Then ∇ and ∇_c satisfy the postulates $(IC0)$, $(IC2)$, $(IC3)$, $(IC5)$, $(IC7)$, $(IC8)$, as well as versions of $(IC1)$, $(IC4)$, $(IC6)$:

$(IC1')$ *If for every $K \in \mathcal{K}$ we have $K \not\vdash \bot$, and $\mu \not\vdash \bot$ then $\nabla^{\mu}(\mathcal{K}) \not\vdash \bot$.* [9]
$(IC4')$ *If $K_1 \not\vdash \bot$, $K_2 \not\vdash \bot$ and $K_1 \vdash \mu$, $K_2 \vdash \mu$ then: $\nabla^{\mu}(\{K_1\} \sqcup \{K_2\}) \wedge K_1 \not\vdash \bot$.*
$(IC6')$ *If $K_1 \wedge K_2 \not\vdash \bot$ then $\nabla(\{K_1\} \sqcup \{K_2\}) \vdash K_1 \wedge K_2$.*

Theorem 3.5 *Let ∇ and ∇_c be defined as in Definition 3.6.*

Then, ∇ and ∇_c satisfy the postulates $(LS1)$–$(LS3)$, $(LS5)$, along with:

$(LS4)'$ *$\alpha \nabla \beta$ is satisfiable iff α is satisfiable and β is satisfiable.*
$(LS8)'$ *If α is satisfiable and β is satisfiable then $\alpha \wedge (\alpha \nabla \beta)$ is satisfiable.*

As well, versions for ∇_c for $(LS4)'$ and $(LS8)'$ also hold.

Postulate $(LS6)$ does not hold here; Example 3 provides a counterexample. As well, the weaker postulate $(LS6)'$ does not hold. Recall that $(LS6)'$ is $(\alpha \nabla \beta_1) \wedge \beta_2$ implies $\alpha \nabla (\beta_1 \wedge \beta_2)$. However, consider the counterexample, derived from the stock-moving example (2):

$$[(p \wedge q) \nabla (\neg p \wedge \neg q)] \ \wedge \ (p \wedge \neg q)$$

---

[9] It is straightforward to obtain $(IC1)$ by essentially ignoring inconsistent belief sets. We remain with the present postulate since it reflects the most natural formulation of projection in our framework.

does not imply

$$(p \wedge q) \ \triangledown \ [(\neg p \wedge \neg q) \wedge (p \wedge \neg q)].$$

Further, postulate $(LS7)$ does not hold here, as Example 6 illustrates. Hence, projection is not an arbitration operator (in the sense of [4]). Neither is the projection operator associative, as the example from the previous subsection, viz. $(((p \vee q)\triangledown\neg p)\triangledown p) \neq (p \vee q)\triangledown(\neg p\triangledown p)$, shows.

Last we have the following results relating projection with merging and revision, respectively:

Theorem 3.6 *Let $\mathcal{K}, \Delta$ and $\triangledown$ be given as in Definitions 3.3 and 3.6 (respectively). Then*

$$\triangledown(\mathcal{K}) \ \subseteq \ \Delta(\mathcal{K}).$$

That is, in binary terms, $\alpha\triangledown\beta \ \subseteq \ \alpha \vartriangle \beta$.

As well, we have the following analogue to Theorem 3.3:

Theorem 3.7 *Let $\dot{+}$ and $\triangledown$ be given as in Definitions 2.2 and 3.6 (respectively).*

Then, $\alpha\dot{+}\beta \ = \ \alpha\triangledown^{\beta}\top$.

3.4 Combining the Approaches

In this section we show an interesting relationship between belief set merging and projection. Consider the belief multiset $\{R\} \sqcup \{K_1, \ldots, K_n\}$. One can define a new type of merging operation in which $\{K_1, \ldots, K_n\}$ are merged (as in Definition 3.3) while simultaneously being projected onto R (as in Definition 3.6). Of course, this is just the merging of $\{R, K_1, \ldots, K_n\}$ (according to Definition 3.2), where instead of taking formulas common to all belief sets, as in Definition 3.3, one just selects those formulas in (the language of) R. This approach then would appear to provide a means of incorporating integrity constraints in belief set merging, in which the belief set R would represent (entailment-based) integrity constraints. However, as will be seen, this enhanced approach adds nothing over what we already have with belief set projection.

Formally we have the following. First, we have the definition, extending Definition 3.2:

Definition 3.7 *Let $B = (\mathcal{K}, R, C)$ be a multi belief change scenario in $\mathcal{L}_{\mathcal{P}}$, where*

$\mathcal{K} = (K_j)_{j \in J}$. *Define* EQ *as a maximal set of equivalences*

$$EQ \subseteq \{p^k \equiv p^l \mid p \in \mathcal{P} \text{ and } k, l \in J\} \cup \{p^j \equiv p \mid p \in \mathcal{P} \text{ and } j \in J\}$$

such that

$$Cn\Big(\bigcup_{j \in J} K_j^j \cup R \cup EQ\Big) \cap (C \cup \{\bot\}) = \emptyset.$$

Then

$$Cn\Big(\bigcup_{j \in J} K_j^j \cup R \cup EQ\Big) \cap \mathcal{L}_{\mathcal{P}}$$

is a consistent merge/project belief change extension *of* B.

If there is no such set EQ *then* B *is* inconsistent *and* $\mathcal{L}_{\mathcal{P}}$ *is defined to be the sole* (inconsistent) symmetric belief change extension *of* B.

We could then go on and define choice and skeptical versions of this operator, as we did for merging and projection. However the following result shows that this is not necessary.

Theorem 3.8 *Let* $B = (\mathcal{K}, R, \emptyset)$ *be a multi belief change scenario in* $\mathcal{L}_{\mathcal{P}}$.

Then E *is a consistent projected belief change extension of* B *(according to Definition 3.5) iff* E *is a consistent merge/project belief change extension of* B *(according to Definition 3.7).*

Consequently the merge/project belief change extensions of B are exactly the projected belief change extensions of B. This in turn means that our purported combining of merging and projection in fact amounts to projection. Expressed differently we have that, in projecting onto R, allowing interactions among the members of $\mathcal{K}$ is in fact irrelevant with respect to the projection.

4 Additional Merging Operators

In this section we outline further types of merging operators. Our results here are less formal than in the previous section, partly because we modify or augment earlier definitions, and partly because the extensions we describe are relatively straightforward. Nonetheless, the operators described here illustrate the range of possibilities that may be covered in the overall approach.

4.1 Rigid merging

More skeptical versions of merging and projection, respectively, can be obtained by introducing some "rigidity" into the definition of merging operators by modifying the format of the sets of equivalences EQ determining belief change extensions.

As regards Definition 3.2, we may replace $EQ \subseteq \{p^k \equiv p^l \mid p \in \mathcal{P} \text{ and } k, l \in J\}$ by

$$EQ \subseteq \{\bigwedge_{k,l \in J}(p^k \equiv p^l) \mid p \in \mathcal{P}\} \tag{1}$$

or even restricted to n-ary conjunctions by

$$EQ \subseteq \{\bigwedge_{k,l \in I}(p^k \equiv p^l) \mid p \in \mathcal{P} \text{ and } I \subseteq J, |I| = n\} .$$

The result however in general is a quite weak operator. Consider for example where we have

$$K_1 = Cn(p \wedge \neg q)\,, \quad K_2 = Cn(q \wedge \neg r)\,, \quad K_3 = Cn(r \wedge \neg p)\,.$$

Then redefining EQ as in (1) we obtain that $\Delta(\{K_1, K_2, K_3\}) = Cn(\top)$. Plausibly however one might expect that $\Delta(\{K_1, K_2, K_3\})$ would contain $p \vee q \vee r$ as well as $\neg p \vee \neg q \vee \neg r$.

More interesting arguably is belief set projection, since the "rigidity" provides us with an additional means for controlling projection. As regards Definition 3.5, we can replace $EQ \subseteq \{p^j \equiv p \mid p \in \mathcal{P} \text{ and } j \in J\}$ with

$$EQ \subseteq \{\bigwedge_{j \in J}(p^j \equiv p) \mid p \in \mathcal{P}\}$$

or even restricted to n-ary conjunctions by

$$EQ \subseteq \{\bigwedge_{j \in I}(p^j \equiv p) \mid p \in \mathcal{P} \text{ and } I \subseteq J, |I| = n\} . \tag{2}$$

As an example, consider four knowledge bases

$$\begin{aligned}
K_1 &\equiv p \wedge q \wedge r \\
K_2 &\equiv p \wedge q \wedge r \\
K_3 &\equiv p \wedge q \wedge \neg r \\
K_4 &\equiv p \wedge \neg q \wedge \neg r\,.
\end{aligned}$$

Now, let us merge these bases by using the type of equivalences in (2) and setting $n = 3$. That is, we only consider conjunctions of three equivalences, like $(q^2 \equiv q) \wedge (q^3 \equiv q) \wedge (q^4 \equiv q)$. Starting from $K_1^1 \cup K_2^2 \cup K_3^3 \cup K_4^4$, we may clearly add all ternary equivalences involving p. Regarding q, however, only the above conjunction can be added, and this is all. As a result, we thus get $p \wedge q$. This is not obtainable with any skeptical merging operator discussed in the previous sections.

Interestingly, the above provides us with a type of majority operator (see (Maj) in Section 2). In fact, adding

$$K_5 \;\equiv\; p \wedge q \wedge r$$

to the four previous knowledge bases triplicates the information in K_1, K_2, or K_5, respectively. Now, merging the five knowledge bases as above gives us in addition to $p \wedge q$ also r, which corresponds to the information contained in K_1, K_2, and K_5. More generally, given m knowledge bases, then merging with sets of equivalences of form (2) where $n > \frac{m}{2}$ satisfies a form of the majority principle. For instance, consider $K_1 = Cn(p)$ and $K_2 = Cn(\neg p)$. While merging $(\{K_1\}^{+m} \sqcup \{K_2\})$ gives $Cn(\top)$ for every $m > 0$ for both belief set merging and projection in Section 3, under the above "majority" operator, we obtain K_1 as the merge of $(\{K_1\}^{+m} \sqcup \{K_2\})$ for $m > 1$.

4.2 Prioritised Merging

Often sources of knowledge are not equally reliable. This is not necessarily an absolute criterion in the sense that one source is generally more reliable than another one, but rather a matter of expertise in the sense that one source is more authoritative on certain subjects, while on others roles may well be interchanged. So for instance, if you want to gather information for an upcoming journey, you might want to prefer (in case of conflict) the weather information from one site and the public transport information from another, although both sites provide information on both topics.

In our setting, this amounts to attributing to sources different priorities on different parts of the alphabet, the idea being that such a part defines the language of a certain subject. For implementing this, we can take advantage of approaches to preference handling in consistency-based reasoning. Among them, let us follow the one of *preferred subtheories* [22] because of its appealing simplicity.

To begin with, consider a family of knowledge bases $(K_j)_{j\in J}$. We express the priorities among these knowledge bases with respect to different subjects by means of a *hierarchy* on the alphabets $\mathcal{P}^J$: A hierarchy associated with a family of knowledge bases $(K_j)_{j\in J}$ is a strict n-ary partition of $\bigcup_{j\in J} \mathcal{P}^j$ for $n > 0$. That is, if

$(P_1, \ldots, P_n)$ is such a hierarchy, we have $\bigcup_{1 \le i \le n} P_i = \bigcup_{j \in J} \mathcal{P}^j$ and $P_i \cap P_j = \emptyset$ for $1 \le i, j \le n$. Intuitively, items in P_i are preferred to those P_j whenever $i < j$. That is, for example, $q^1 \in P_5$ and $q^2 \in P_7$ reflects the idea that the contents of knowledge base K_5 regarding q is considered more reliable than that in K_7.

Now, we can use the information in a hierarchy to guide the formation of a maximal set of equivalences. We do this in the context of projected merging, since, as above, this type of merging is more easily parameterizable than symmetric merging.

Definition 4.1 *Let $H = (P_1, \ldots, P_n)$ be a hierarchy associated with the family of knowledge bases $(K_j)_{j \in J}$.*

Define EQ as an H-maximal set of equivalences, if $EQ = \bigcup_{1 \le i \le n} EQ_i$ and for all k such that $1 \le k \le n$ we have that $\bigcup_{1 \le i \le k} EQ_i$ is a maximal set of equivalences

$$\bigcup_{1 \le i \le k} EQ_i \subseteq \{p^j \equiv p \mid p \in \mathcal{P}, j \in J, \textit{ and } p^j \in \bigcup_{1 \le i \le k} P^i\}$$

such that

$$Cn\Big(\bigcup_{j \in J} K_j^j \cup R \cup \bigcup_{1 \le i \le k} EQ_i\Big) \cap (C \cup \{\bot\}) = \emptyset.$$

Then

$$Cn\Big(\bigcup_{j \in J} K_j^j \cup R \cup EQ\Big) \cap \mathcal{L}_{\mathcal{P}}$$

is a consistent H-projected belief change extension *of B.*

If there is no such set EQ then B is inconsistent *and $\mathcal{L}_{\mathcal{P}}$ is defined to be the sole* (inconsistent) H-projected belief change extension *of B.*

The definition of a merging operation is analogous to the previous ones.

As an example, let us reconsider the knowledge bases from Example 2:

$$K_1 \equiv p \wedge q \wedge r \wedge s \quad \text{and} \quad K_2 \equiv \neg p \wedge \neg q \wedge \neg r \wedge \neg s.$$

Suppose that the first analyst is more competent on stocks p and q, while the second is more qualified for stocks r and s. This can be modeled through the following hierarchy:

$$H = (\{p^1, q^1, r^2, s^2\}, \{p^2, q^2, r^1, s^1\}) .$$

We get a single H-projected belief change extension being equivalent to $p \wedge q \wedge \neg r \wedge \neg s$.

Interestingly, this hierarchical form of merging can be put in correspondence to revision, and this in a different way than done in the previous section. Assume that we are given two consistent formulas,[10] K_1 and K_2 over alphabet $\mathcal{P}$. Then, the revision of K_1 by K_2, that is, $K_1 \dot{+} K_2$, corresponds to the H-projected merge of K_1 and K_2 with respect to hierarchy $H = (\{p^1 \mid p \in \mathcal{P}\}, \{p^2 \mid p \in \mathcal{P}\})$.

Sometimes it is desirable that knowledge bases remain *neutral* with respect to certain propositions. Consider where we have knowledge bases K_1 and K_2 and alphabet $\{p, q, r\}$. K_1 is trusted more wrt p, and K_2 is trusted more wrt q, but neither is preferred wrt r. Having the hierarchy $(\{p_1, q_2, r_1, r_2\}, \{p_2, q_1\})$, for instance, doesn't capture this since it ranks r_1, r_2 wrt p_2, q_1. Arguably the partition should be $(\{p_1, q_2\}, \{p_2, q_1\})$, which less constrains the possible EQ sets. This can be accommodated as follows. A hierarchy is only defined on a certain subset of $\mathcal{P}^J$. Let N contain the non-ranked propositions, that is, $\bigcup_{j \in J} \mathcal{P}^j = N \cup \bigcup_{1 \leq i \leq n} P_i$. Then, it is sufficient to replace $p^j \in \bigcup_{1 \leq i \leq k} P^i$ in Definition 4.1 by $p^j \in \bigcup_{1 \leq i \leq k} P^i \cup N$ in order to obtain the desired result.

Finally, let us mention that our way of ranking can also be used to put priorities on general formulas. If K_1 is to be trusted over K_2 wrt ϕ, then one needs just introduce a new atom p_ϕ, along with assertion (or new integrity constraint) $(p_\phi \equiv \phi)$, and then assert that K_1 is to be trusted over K_2 wrt p_ϕ.

5 Complexity

In [23], we analysed the computational complexity of reasoning from belief change scenarios. Specifically, we addressed the following basic reasoning tasks:

Theorem 5.1 ([23])

(1) Deciding whether a belief change scenario B has a consistent belief change extension is NP-complete;

(2) Given a belief change scenario B and formula ϕ, deciding whether ϕ is contained in at least one consistent belief change extension of B is Σ_P^2-complete; and

(3) Given a belief change scenario B and formula ϕ, deciding whether ϕ is contained in all consistent belief change extensions of B is Π_P^2-complete.

Clearly, the variants of these decision problems for merging and projection fall in the same complexity class and in fact follow as corollaries of the above result. This then illustrates an advantage of formulating belief change operations within a uniform framework: essentially, properties of the basic framework can be investigated

[10] or indeed sets of formulas.

in a general form; properties of specific operators (or combinations of operators) are then easily derivable as secondary results.

6 Discussion

We have presented a general consistency-based framework for specifying belief set merging operators. Two major approaches for merging belief sets were developed. In the first approach, the intuition is that for merging belief sets, common information is in a sense "pooled". This approach then seems to conform to the commonsense notion of merging of knowledge, in which belief sets are joined to produce a single belief set retaining as much as possible of the contents of the original belief sets. A characteristic of this operation is that sentences common to the original belief sets are in the merged belief set.

In the second approach, belief sets are projected onto another belief set. That is, the sets we wish to merge are used to augment the knowledge of another (possibly empty) belief set. This second approach appears to differ from others that have appeared in the literature. It is strictly weaker than the first; however this weakness is not a disadvantage, since, among other things, it avoids the possible difficulty illustrated in Example 2. This second approach has something of the flavour of both belief revision and update. With respect to belief revision, projection can be viewed as a process whereby several belief sets are simultaneously revised with respect to another. With respect to belief update, semantically, individual models of a belief set are independently updated. Hence projection is like update, but where the "granularity" of the operation is at the level of belief sets rather than models. Thus projection can be regarded as an operator lying intermediate between belief revision and update.

The role of integrity constraints was examined in these approaches. As well, we also more briefly considered variant merging operators, including a prioritised approach to merging, wherein different knowledge sources could within them have (relative) varying levels of reliability.

Acknowledgements

We would like to express our great thanks to Jérôme Lang for many helpful suggestions and fruitful discussions on earlier drafts of this paper. We also thank the referees of [24,21] for their detailed and helpful comments.

A Proofs

Proof 3.1

The statement of the theorem has $R = C = \emptyset$, which corresponds to $\mu = \top$ in Definition 2.4. With $\mu = \top$ we obtain the postulates of [1], where integrity constraints are not addressed. We remain with Definition 2.4 for uniformity with our consideration of the project operator, following.

$(IC0)$, $(IC1')$, $(IC2)$, and $(IC3)$ are all obvious from the definition of merge.

For $(IC4')$, assume that $K \nvdash \bot$ for every $K \in \mathcal{K}$, but that there is $K_i \in \mathcal{K}$ where $\Delta(\mathcal{K}) \cup K_i \vdash \bot$.

From Definition 3.2 we obtain that

$$Cn\Big(\bigcup_{j \in J} K_j^j \cup EQ\Big) \cup K_i^i \vdash \bot$$

for every set EQ satisfying the terms of the definition.

However, trivially we also have that $K_i^i \subseteq Cn\Big(\bigcup_{j \in J} K_j^j \cup EQ\Big)$ for every set EQ satisfying the terms of the definition, and so $\Big(\bigcup_{j \in J} K_j^j\Big) \cup EQ \vdash \bot$.

But this contradicts $(IC1')$, and so there is no $K_i \in \mathcal{K}$ where $\Delta(\mathcal{K}) \cup K_i \vdash \bot$, establishing what was to be shown.

For $(IC5)$, let $\mathcal{K}_1 = (K_1, \ldots, K_n)$ and $\mathcal{K}_2 = (K_{n+1}, \ldots, K_m)$.

If $\Delta(\mathcal{K}_1) \wedge \Delta(\mathcal{K}_2) \vdash \bot$ then the result is immediate; hence assume that $\Delta(\mathcal{K}_1) \wedge \Delta(\mathcal{K}_2) \nvdash \bot$.

Let E_1 be a symmetric belief change extension of $(\mathcal{K}_1, \emptyset, \emptyset)$ over $\mathcal{P}^{\{1..n\}}$ with corresponding set of equivalences EQ_1, and let E_2 be a symmetric belief change extension of $(\mathcal{K}_2, \emptyset, \emptyset)$ over $\mathcal{P}^{\{n+1..m\}}$ with corresponding set of equivalences EQ_2.

Clearly $EQ_1 \cup EQ_2$ can be extended to a set of equivalences over $\mathcal{P}^{\{1..m\}}$ for $(K_1, \ldots, K_m)$, from which our result follows.

That we do not obtain $(IC6)$ follows from [5, Theorem 3.3]: a merging operator that satisfies $(IC2)$, $(IC4)$, and $(IC6)$ cannot satisfy majority independence. Since Δ satisfies $(IC2)$, $(IC4)$, and majority independence it cannot thereby satisfy $(IC6)$.

$(IC7)$ and $(IC8)$ are trivial here. ■

Proof 3.2

$(LS1)$, $(LS4)'$, and $(LS5)$ are obvious.

For $(LS2)$, if $\vdash \neg(\alpha \wedge \beta)$ then the result is immediate.

If $\not\vdash \neg(\alpha \wedge \beta)$ then there is a unique set of equivalences determining a single belief change extension to $(\{\alpha\}, \{\beta\}, \emptyset)$ given by $EQ = \{p \equiv p' \mid p \in \mathcal{P}\}$. It follows in this case that $(\alpha \wedge \beta) \equiv (\alpha \vartriangle \beta)$. This also serves to show $(LS3)$.

For $(LS6)'$, we need to show that $(\alpha \vartriangle \beta_1) \wedge \beta_2$ implies $\alpha \vartriangle (\beta_1 \wedge \beta_2)$, or that if

$$\bigcap_{i \in I} Cn(\{\alpha'\} \cup \{\beta_1 \wedge \beta_2\} \cup EQ_i) \vdash \phi \wedge \phi'$$

then

$$\left(\bigcap_{i \in I} Cn(\{\alpha'\} \cup \{\beta_1\} \cup EQ_i)\right) \cup \{\beta_2\} \vdash \phi \wedge \phi'.$$

The proof is the same as that for $(K\dot{+}7)$ in [2, Theorem 4.2].

For $(LS7)$, we show that if $\vdash (\alpha \vee \beta) \supset \phi$ for arbitrary ϕ, then $\vdash (\alpha \vartriangle \beta) \supset \phi$. If $\vdash (\alpha \vee \beta) \supset \phi$ then $\vdash \alpha \supset \phi$ and $\vdash \beta \supset \phi$. But we also have that $\vdash \alpha' \supset \phi'$ and $\vdash \beta' \supset \phi'$ and so by the definition of $\vartriangle$ we get that $\phi \in \alpha \vartriangle \beta$ and so $\vdash \alpha \vartriangle \beta \supset \phi$.

$(LS8)$ follows from the observation that if $\alpha \not\vdash \bot$ and $\beta \not\vdash \bot$ then it is an easy consequence of the definition of $\vartriangle$ that $\alpha \vartriangle \beta \not\vdash \neg\alpha$, $\alpha \vartriangle \beta \not\vdash \neg\beta$. (See as well the proof of $(IC4')$ in the preceding proof.)

For $(LS6c)'$, assume that EQ determines symmetric belief change extension of $(\{\alpha\}, \{\beta_1 \vee \beta_2\}, \emptyset)$, and let c be the function that selects this belief change extension. From [2, Lemma A.1] we have that EQ determines symmetric belief change extension of $(\{\alpha\}, \{\beta_1\}, \emptyset)$ or of $(\{\alpha\}, \{\beta_2\}, \emptyset)$. Assume without loss of generality that EQ determines symmetric belief change extension of $(\{\alpha\}, \{\beta_1\}, \emptyset)$, and let c' be the function that selects this belief change extension.

We have that $\alpha \vartriangle_c \beta_1$ is $\{\phi \mid \{\alpha'\} \cup \{\beta_1\} \cup EQ \vdash \phi \wedge \phi'\}$ and $\alpha \vartriangle_c (\beta_1 \vee \beta_2)$ is $\{\phi \mid \{\alpha'\} \cup \{\beta_1 \vee \beta_2\} \cup EQ \vdash \phi \wedge \phi'\}$. Consequently we have that $\alpha \vartriangle_c \beta_1$ implies $\alpha \vartriangle_c (\beta_1 \vee \beta_2)$. ■

Proof 3.3

We have that $\phi \in \alpha \vartriangle \beta$ iff for every set of equivalences EQ determining a belief change extension that $\phi, \phi' \in Cn(\alpha' \cup \beta \cup EQ)$, where $\phi \in \mathcal{L}_{\mathcal{P}}$, and $\phi' \in \mathcal{L}_{\mathcal{P}'}$.

But $\phi \in Cn(\alpha' \cup \beta \cup EQ)$, $\phi \in \mathcal{L}_{\mathcal{P}}$ for every such EQ iff $\phi \in \alpha \dot{+} \beta$.

And:

$\phi' \in Cn(\alpha' \cup \beta \cup EQ)$ where $\phi' \in \mathcal{L}_{\mathcal{P}'}$ for every such set EQ

iff $\phi \in Cn(\alpha \cup \beta' \cup EQ)$ where $\phi \in \mathcal{L}_{\mathcal{P}}$ for every such set EQ

iff $\phi \in \beta \dot{+} \alpha$.

Consequently we have $\phi \in \alpha \bigtriangleup \beta$ iff $\phi \in \alpha \dot{+} \beta$ and $\phi \in \beta \dot{+} \alpha$; thus $\alpha \bigtriangleup \beta \equiv \alpha \dot{+} \beta \cap \beta \dot{+} \alpha$.

■

Proof 3.4

$(IC0)$, $(IC1')$, $(IC2)$, and $(IC3)$ are all obvious from the definition of project. (For $(IC2)$ we would have $EQ = \{p^j \equiv p \mid p \in \mathcal{P}$ and $j \in J\}$ from which the result follows.)

For $(IC4')$, since $K_1 \vdash \mu$ we have that there is a maximal set EQ defined using the language $\mathcal{P}^1$ such that $K_1^1 \cup \{\mu\} \cup EQ$ is consistent. Clearly $K_1^1 \cup K_2^2 \cup \{\mu\} \cup EQ$ is consistent. Consequently we can extend EQ to a maximal set EQ', according to Definition 3.5 over language $\mathcal{P}^1 \cup \mathcal{P}^2$ such that $K_1^1 \cup K_2^2 \cup \{\mu\} \cup EQ'$ is consistent, from which our result obtains.

$(IC5)$ is the same as in Theorem 3.1, but allowing $R \neq \emptyset$. $(IC6)$ fails for the same reason as in Theorem 3.1.

For $(IC7)$, assume that $\nabla^{\mu_1}(\mathcal{K}) \wedge \mu_2 \nvdash \bot$ (otherwise our result holds trivially).

We have for every choice of EQ satisfying Definition 3.5 that

$$
\begin{aligned}
\nabla_c^{\mu_1}(\mathcal{K}) \wedge \mu_2 &= Cn\Big((Cn\Big(\bigcup\nolimits_{j\in J} K_j^j \cup \{\mu_1\} \cup EQ\Big) \cap \mathcal{L}_{\mathcal{P}}) \cup \{\mu_2\}\Big) \\
&= Cn\Big((Cn\Big(\bigcup\nolimits_{j\in J} K_j^j \cup \{\mu_1\} \cup EQ \cup \{\mu_2\}\Big) \cap \mathcal{L}_{\mathcal{P}})\Big) \\
&= Cn\Big(\bigcup\nolimits_{j\in J} K_j^j \cup \{\mu_1, \mu_2\} \cup EQ\Big) \cap \mathcal{L}_{\mathcal{P}} \\
&= \nabla_c^{\mu_1 \wedge \mu_2}(\mathcal{K})
\end{aligned}
$$

$(IC8)$ holds, given the consistency condition, by virtue of the fact that we have equalities in the preceding.

■

Proof 3.5

$(LS1)$–$(LS3)$, $(LS4)'$, and $(LS5)$ are obvious. ($(LS2)$ follows analogously to the proof for merging.)

For $(LS6)$, the counterexample is (writing conjunction by juxtaposition and giving conjunction higher precedence than disjunction):

α is $pqrs$, β_1 is $\neg p \neg q \vee \neg r$, β_1 is $\neg q \vee \neg s$.

$\alpha \triangledown (\beta_1 \vee \beta_2)$ is $pqrs \triangledown (\neg q \vee \neg r \vee \neg s)$ which is $p \wedge (qrs \vee \neg qrs \vee q \neg rs \vee qr \neg s)$.

$\alpha \triangledown \beta_1$ is $pqrs \triangledown (\neg p \neg q \vee \neg r)$ which is $s \wedge (pqr \vee \neg pqr \vee p \neg qr \vee pq \neg r \vee \neg p \neg qr)$.

$\alpha \triangledown \beta_2$ is $pqrs \triangledown (\neg q \vee \neg s)$ which is $pr \wedge (qs \vee \neg qs \vee q \neg s)$

$(LS8)'$ follows from the observation that if $\alpha \not\vdash \bot$ and $\beta \not\vdash \bot$ then it is an easy consequence of the definition of $\triangledown$ that $\alpha \triangledown \beta \not\vdash \neg \alpha$. That is, one can choose an initial EQ set by: $EQ = \{p \equiv p^1 \mid p \in \mathcal{P}\}$. The satisfiability of α guarantees that $\alpha \cup EQ$ is satisfiable. EQ can subsequently be extended to a maximum set satisfying Definition 3.5, from which our result follows. ∎

Proof 3.6

Let $B = (\mathcal{K}, \emptyset, \emptyset)$ be a multi belief change scenario. Toward showing that $\nabla(\mathcal{K}) \subseteq \Delta(\mathcal{K})$, choose $\alpha \in \mathcal{L}_{\mathcal{P}}$ such that $\alpha \notin \Delta(\mathcal{K})$. We show that $\alpha \notin \nabla(\mathcal{K})$, as follows.

Since $\alpha \notin \Delta(\mathcal{K})$, by Definition 3.3 there is some maximal set of equivalences EQ and an index i such that $Cn\Big(\bigcup_{j \in J} K_j^j \cup EQ\Big) \cap \{\bot\} = \emptyset$ and $\alpha^i \notin Cn\Big(\bigcup_{j \in J} K_j^j \cup EQ\Big)$.

Define

$$
\begin{aligned}
EQ' &= EQ \cup \{p^i \equiv p \mid p \in \mathcal{P}\} \\
EQ^* &= EQ' \cup \{p^j \equiv p \mid EQ' \vdash p^j \equiv p \text{ where } p \in \mathcal{P} \text{ and } j \in J\} \\
EQ^+ &= EQ^* \setminus \{p^j \equiv p^k \mid p^j \equiv p^k \in EQ, p \in \mathcal{P} \text{ and } j, k \in J\}
\end{aligned}
$$

Clearly

$$
\begin{aligned}
Cn\Big(\textstyle\bigcup_{j \in J} K_j^j \cup EQ'\Big) \cap \{\bot\} &= \emptyset \qquad \text{and} \\
Cn\Big(\textstyle\bigcup_{j \in J} K_j^j \cup EQ^*\Big) \cap \{\bot\} &= \emptyset.
\end{aligned}
$$

As well, we have $Cn(EQ^*) = Cn(EQ^+)$.

We have that EQ^+ satisfies the conditions for a maximal set of equivalences according to Definition 3.6. The argument is as follows:

> We had originally that $Cn\Big(\bigcup_{j\in J} K_j^j \cup EQ\Big) \cap \{\bot\} = \emptyset$ is a maximal set of equivalences according to Definition 3.3.
>
> As mentioned, $Cn\Big(\bigcup_{j\in J} K_j^j \cup EQ'\Big) \cap \{\bot\} = \emptyset$ (since the added equivalences in EQ' involve a new language, viz. $\mathcal{P}$), and
>
> $Cn\Big(\bigcup_{j\in J} K_j^j \cup EQ^*\Big) \cap \{\bot\} = \emptyset$ (since EQ^* simply includes derivable equivalences).
>
> From $Cn(EQ^*) = Cn(EQ^+)$ we get that $Cn(EQ^+)$ determines a maximal set of equivalences according to Definition 3.6.

Now let M be a model of $\bigcup_{j\in J} K_j^j \cup EQ$ such that $M \not\models \alpha^i$, and let M' be its extension to a model of $\bigcup_{j\in J} K_j^j \cup EQ'$.

It follows from the preceding that M' is a model of $\bigcup_{j\in J} K_j^j \cup EQ^+$.

But this means that $M' \not\models \alpha$ given the equivalences in EQ', which are retained in EQ^+.

Thus we have that M' is a model of $\bigcup_{j\in J} K_j^j \cup EQ^+$, $M' \not\models \alpha$, and EQ^+ determines a maximal set of equivalences according to Definition 3.6.

Consequently $\alpha \notin \nabla(\mathcal{K})$, which was to be shown. ■

Proof 3.7

This is an easy consequence of Definitions 2.2 and 3.6.

■

Proof 3.8

Note that Definition 3.7 extends Definition 3.5 by the addition of the term

$$\{p^k \equiv p^l \mid p \in \mathcal{P} \text{ and } k, l \in J\} \tag{A.1}$$

in the specification of EQ.

Let EQ be a set of equivalences according to the given conditions in Definition 3.5. For convenience, let

$$\Gamma = Cn\Big(\bigcup_{j\in J} K_j^j \cup R \cup EQ\Big).$$

Assume that $\Gamma \not\vdash \bot$; otherwise our result follows trivially.

We prove the result by showing that for every $p^k \equiv p^l$, $k, l \in J$, either $\Gamma \vdash p^k \equiv p^l$ or $\Gamma \vdash \neg(p^k \equiv p^l)$. That is, the extra term (A.1) in the definition of EQ in Definition 3.7 is in fact redundant.

Let EQ be a set of equivalences according to Definition 3.5. Let $p \in \mathcal{P}$. For $j, k \in J$ we have the following possibilities:

(1) $p^j \equiv p \in EQ$ and $p^k \equiv p \in EQ$.
We have that $EQ \vdash p^j \equiv p^k$ and so from the monotonicity of classical logic we obtain $\Gamma \vdash p^j \equiv p^k$.

(2) $p^j \equiv p \in EQ$ and $p^k \equiv p \notin EQ$.
By assumption we have that $\Gamma \nvdash \perp$ and from the maximality of EQ we have that $\Gamma \cup \{p^k \equiv p\} \vdash \perp$. Hence $\Gamma \vdash \neg(p^k \equiv p)$, and since by assumption we have that $\Gamma \vdash p^j \equiv p$ we obtain that $\Gamma \vdash \neg(p^k \equiv p^j)$.

(3) $p^j \equiv p \notin EQ$ and $p^k \equiv p \in EQ$.
This is the same as case 2 above.

(4) $p^j \equiv p \notin EQ$ and $p^k \equiv p \notin EQ$.
From the maximality of EQ we have that $\Gamma \cup \{p^j \equiv p\} \vdash \perp$ and $\Gamma \cup \{p^k \equiv p\} \vdash \perp$. Thus $\Gamma \vdash \neg(p^j \equiv p)$ and $\Gamma \vdash \neg(p^k \equiv p)$ from which we obtain that $\Gamma \vdash p^k \equiv p^j$.

This shows that for every $p^k \equiv p^l$, $k, l \in J$, either $\Gamma \vdash p^k \equiv p^l$ or $\Gamma \vdash \neg(p^k \equiv p^l)$, which was to be shown. ■

References

[1] S. Konieczny, R. Pino Pérez, On the logic of merging, in: A. G. Cohn, L. Schubert, S. C. Shapiro (Eds.), KR'98: Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, San Francisco, California, 1998, pp. 488–498.

[2] J. Delgrande, T. Schaub, A consistency-based approach for belief change, Artificial Intelligence 151 (1-2) (2003) 1–41.

[3] P. Gärdenfors, Knowledge in Flux: Modelling the Dynamics of Epistemic States, The MIT Press, Cambridge, MA, 1988.

[4] P. Liberatore, M. Schaerf, Arbitration (or how to merge knowledge bases), IEEE Transactions on Knowledge and Data Engineering 10 (1) (1998) 76–90.

[5] S. Konieczny, R. Pino Pérez, Merging information under constraints: A logical framework, Journal of Logic and Computation 12 (5) (2002) 773–808.

[6] M. Dalal, Investigations into theory of knowledge base revision., in: Proceedings of the AAAI National Conference on Artificial Intelligence, St. Paul, Minnesota, 1988, pp. 449–479.

[7] C. Baral, S. Kraus, J. Minker, V. Subrahmanian, Combining multiple knowledge bases consisting of first order theories, Computational Intelligence 8 (1) (1992) 45–71.

[8] P. Revesz, On the semantics of theory change: Arbitration between old and new information, in: C. Beeri (Ed.), Proceedings of the Twelth ACM Symposium on Principles of Database Systems, Washington D.C., 1993, pp. 71–82.

[9] S. Konieczny, On the difference between merging knowledge bases and combining them, in: A. G. Cohn, F. Giunchiglia, B. Selman (Eds.), KR2000: Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, San Francisco, 2000, pp. 135–144.

[10] P. Liberatore, M. Schaerf, Reducing belief revision to circumscription (and vice versa), Artificial Intelligence 93 (1–2) (1997) 261–296.

[11] J. Lin, A. Mendelzon, Knowledge base merging by majority, in: R. Pareschi, B. Fronhöfer (Eds.), Dynamic Worlds: From the Frame Problem to Knowledge Management, Kluwer, 1999.

[12] S. Konieczny, J. Lang, P. Marquis, Distance-based merging: a general framework and some complexity results, in: D. Fensel, F. Giunchiglia, D. McGuiness, M. Williams (Eds.), Proceedings of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, San Francisco, 2002, pp. 97–108.

[13] T. Meyer, On the semantics of combination operations, Journal of Applied NonClassical Logics 11 (1-2) (2001) 59–84.

[14] W. Spohn, Ordinal conditional functions: A dynamic theory of epistemic states, in: W. Harper, B. Skyrms (Eds.), Causation in Decision, Belief Change, and Statistics, Vol. II, Kluwer Academic Publishers, 1988, pp. 105–134.

[15] R. Booth, Social contraction and belief negotiation, in: D. Fensel, F. Giunchiglia, D. McGuiness, M. Williams (Eds.), Proceedings of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, San Francisco, 2002, pp. 375–384.

[16] S. Benferhat, D. Dubois, S. Kaci, H. Prade, Possibilistic merging and distance-based fusion of prop ositional information, Annals of Mathematics and Artificial Intelligence 34 (1-3) (2003) 217–252.

[17] R. Reiter, Towards a logical reconstruction of relational database theory, in: M. Brodie, J. Mylopoulos, J. Schmidt (Eds.), On Conceptual Modelling, Springer-Verlag, 1984, pp. 191–233.

[18] R. Kowalski, Logic for data description, in: H. Gallaire, J. Minker (Eds.), Logic and Data Bases, Plenum Press, 1978, pp. 77–103.

[19] F. Sadri, R. Kowalski, A theorem-proving approach to database integrity, in: J. Minker (Ed.), Foundations of Deductive Databases and Logic Programming, Morgan Kaufmann Publishers, 1987, Ch. 9, pp. 313–362.

[20] H. Katsuno, A. Mendelzon, Propositional knowledge base revision and minimal change, Artificial Intelligence 52 (3) (1991) 263–294.

[21] J. Delgrande, T. Schaub, Consistency-based approaches to merging knowledge bases, in: J. Alferes, J. Leite (Eds.), Proceedings of the Ninth European Conference on Logics in Artificial Intelligence (JELIA'04), Lecture Notes in Artificial Intelligence, 2004, to appear.

[22] G. Brewka, Preferred subtheories: An extended logical framework for default reasoning, in: Proceedings of the International Joint Conference on Artificial Intelligence, 1989, pp. 1043–1048.

[23] J. Delgrande, T. Schaub, H. Tompits, S. Woltran, On computing solutions to belief change scenarios, in: S. Benferhat, P. Besnard (Eds.), Proceedings of the Sixth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2001), Vol. 2143 of Lecture Notes in Artificial Intelligence, Springer Verlag, Toulouse, Fr., 2001, pp. 510–521.

[24] J. Delgrande, T. Schaub, Consistency-based approaches to merging knowledge bases, in: J. Delgrande, T. Schaub (Eds.), Proceedings of the Tenth International Workshop on Non-Monotonic Reasoning (NMR 2004), 2004, pp. 126–133.