

Two Approaches to Merging Knowledge Bases

James P. Delgrande¹ and Torsten Schaub^{2*}

¹ School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada V5A 1S6,
jim@cs.sfu.ca

² Institut für Informatik, Universität Potsdam, D-14415 Potsdam, Germany,
torsten@cs.uni-potsdam.de

Abstract. We present a framework for expressing different merging operators for belief sets. This framework is a generalisation of our earlier work concerning consistency-based belief revision and contraction. Two distinct merging operators are identified: in the first approach, belief sources are consistently combined so that the result of merging knowledge bases K_1, \dots, K_n is a maximal consistent (if possible) set of formulas comprising the joint knowledge of the knowledge bases. This approach then accords to 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, \dots, K_n are "projected" onto another (in the simplest case the trivially true knowledge base). In both cases, we consider the incorporation of entailment-based and consistency-based integrity constraints. Properties of these operators are investigated, primarily by comparing their properties with postulates that have been identified previously in the literature. As well, the interrelationships between these approaches and belief revision is given.

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 [11]) *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

* Affiliated with the School of Computing Science at Simon Fraser University, Canada.

paper, we develop a specific framework for specifying 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 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 one of the knowledge bases contains correct information, but it is not known which.

In the second approach, knowledge bases are projected onto a separate knowledge base (which in the simplest case would consist solely of the trivially true knowledge base). 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; thus 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 incorporation of entailment-based and consistency-based integrity constraints with the merge operator. 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 our approaches. We conclude with a discussion. Proofs are omitted due to space limitations.

2 Background

2.1 Consistency-Based Belief Revision

This subsection summarises our earlier work in [5]. Throughout this paper, we deal with propositional languages and use the logical symbols \top , \perp , \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 , \dots . *Knowledge bases* are identified with deductively-closed sets of formulas, or *belief sets*, and are denoted K , K_1 , \dots .³ Thus $K = Ch(K)$, where $Ch(\cdot)$ 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

³ We note that while we deal solely with belief sets in this paper, our definitions work for arbitrary sets of formulas, and provide the basis for a finite representation of these operators.

$p \in \mathcal{P}$ by the corresponding proposition $p' \in \mathcal{P}'$ (so implicitly there is an isomorphism between \mathcal{P} and \mathcal{P}'). This is defined analogously for sets of formulas.

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.

Definition 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 will 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 will use a selection function to select a specific consistent belief change extension.⁴

Definition 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 (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 .

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

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 $\{\neg p \equiv q\} \dot{+} \neg q = Cn(p \wedge \neg q)$ by $\{p \equiv p', q \equiv q'\}$ from $\neg p' \equiv q'$ and $\neg q$. 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 determined by $\{p \equiv p', q \equiv q'\}$, while the latter relies on two such sets, viz. $\{p \equiv p'\}$ and $\{q \equiv q'\}$.

⁴ This use of selection functions is slightly different from that in the AGM approach.

Definition 1 also leads to a natural and general treatment of both consistency-based and entailment-based integrity constraints; see [5] for details.

2.2 Belief Merging

Konieczny and Pino Peréz [9] consider the problem of merging possibly contradictory belief bases. To this end, they consider finite multisets of the form $\Psi = \{K_1, \dots, K_n\}$ 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 . Multiset union is denoted \sqcup , wherein for example $\{\phi\} \sqcup \{\phi\} = \{\phi, \phi\}$. Following [9], we use⁵ $\Delta^\mu(\Psi)$ to denote the result of merging the multi-set Ψ of belief bases given the entailment-based integrity constraint expressed by μ . They provide the following set of postulates:

Definition 3 ([9]). *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 $\mu \not\vdash \perp$ then $\Delta^\mu(\Psi) \not\vdash \perp$.*
- (IC2) *If $\bigwedge \Psi \not\vdash \neg\mu$ 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: $\Delta^\mu(\phi \sqcup \phi') \wedge \phi \not\vdash \perp$ implies $\Delta^\mu(\phi \sqcup \phi') \wedge \phi' \not\vdash \perp$.*
- (IC5) $\Delta^\mu(\Psi_1) \wedge \Delta^\mu(\Psi_2) \vdash \Delta^\mu(\Psi_1 \sqcup \Psi_2)$.
- (IC6) *If $\Delta^\mu(\Psi_1) \wedge \Delta^\mu(\Psi_2) \not\vdash \perp$ 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 $\Delta^{\mu_1}(\Psi) \wedge \mu_2 \not\vdash \perp$ 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 [11].

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 [9] for an explanation.

⁵ [11] write $\Delta_\mu(\Psi)$ where we have $\Delta^\mu(\Psi)$.

(*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 $\Delta^{\mu_1 \vee \mu_2}(\phi_1 \sqcup \phi_2) \equiv \Delta^{\mu_1}(\phi_1)$.

[11] 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 [4].

Liberatore and Schaerf [13] consider merging two belief bases and propose the following postulate set to characterise a merge operator that they call an *arbitration* operator and that [9] call a *commutative revision operator*. Like [9] they restrict their attention to propositional languages over a finite set of atoms.

(*LS1*) $\vdash \alpha \Delta \beta \equiv \beta \Delta \alpha$.

(*LS2*) $\vdash \alpha \wedge \beta \supset \alpha \Delta \beta$.

(*LS3*) If $\alpha \wedge \beta$ is satisfiable then $\vdash \alpha \Delta \beta \supset \alpha \wedge \beta$.

(*LS4*) $\alpha \Delta \beta$ 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 $\vdash \alpha_1 \Delta \beta_1 \equiv \alpha_2 \Delta \beta_2$.

(*LS6*) $\alpha \Delta (\beta_1 \vee \beta_2) = \begin{cases} \alpha \Delta \beta_1 & \text{or} \\ \alpha \Delta \beta_2 & \text{or} \\ (\alpha \Delta \beta_1) \vee (\alpha \Delta \beta_2) \end{cases}$

(*LS7*) $\vdash (\alpha \Delta \beta) \supset (\alpha \vee \beta)$.

(*LS8*) If α is satisfiable then $\alpha \wedge (\alpha \Delta \beta)$ is satisfiable.

Earlier work on merging operators includes [1] and [17]. The former proposes various theory merging operators based on the selection of maximum consistent subsets in the union of the belief bases; see [10] for a pertinent discussion. The latter proposes an "arbitration" operator that satisfies a subset of the Liberatore and Schaerf postulates; see [12] for a discussion. [14] first identified and addressed the majority merge operator. [8] 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, [15] proposed a general approach to formulating merging functions, based on ordinal conditional functions [19]. 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. [3] 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 [2].

3 Consistency-Based Approaches to Belief Set Merging

In this section we modify the framework given by Definition 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, 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. In the second approach, knowledge sources are projected onto a separate knowledge source (which in the simplest case could consist solely of \top). That is, the sources we wish to merge are used to augment the knowledge of another source.

3.1 Multi belief change scenarios

A *multi belief change scenario* 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 and C are sets of formulas in $\mathcal{L}_{\mathcal{P}}$. 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 negations of formulas in C will not be contained in the result. While R is intended to represent a set of entailment-based integrity constraints [16], 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], 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.

Definition 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

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

Then

$$\left\{ \alpha \mid \{\alpha^j \mid j \in J\} \subseteq Cn\left(\bigcup_{j \in J} K_j^j \cup R^J \cup EQ\right) \right\}$$

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 .

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. Similarly, the formulas C^J ensure that the formulas in C will not be in the result.

Definition 5 (Merging). Let \mathcal{K} be a family of sets of formulas in $\mathcal{L}_{\mathcal{P}}$ and R and C be finite 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}, R, C) .

Then, we define

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

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, $\Delta^{R,C}(\{K_1, K_2\})$ is written as $K_1 \Delta^{R,C} K_2$. Also, given two formulas α, β along with $R = C = \emptyset$, we just write $\alpha \Delta \beta$. For conformity with the notation used in Definition 3 [9], we write $\alpha \Delta^\mu \beta$ if $R = \{\mu\}$ and $C = \emptyset$.

Example 1. $(p \wedge q \wedge r) \Delta (p \wedge \neg q \wedge s)$ 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 $K_1 \Delta K_2$ yields $EQ = \emptyset$ and in fact

$$K_1 \Delta K_2 = Cn(\{(p \wedge q \wedge r \wedge s) \vee (\neg p \wedge \neg q \wedge \neg r \wedge \neg s)\})$$

This example is introduced and discussed in [11]; 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, one forecast will be believed to hold in its entirety, or the other will. As [11] 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 forecast 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.

Theorem 1. Let Δ^μ and Δ_c^μ be defined as in Definition 5.

Then Δ^μ and Δ_c^μ satisfy the postulates (IC0), (IC2) – (IC5), (IC7) – (IC8), as well as the weaker version of (IC1).⁶

(IC1') If $K \not\vdash \neg\mu$ for every $K \in \Psi$ and $\mu \not\vdash \perp$ then $\Delta^\mu(\Psi) \not\vdash \perp$.

A counterexample to (IC6) is given by $\Psi_1 = \{Cn(p), Cn(\neg p)\}$, $\Psi_2 = \{Cn(p)\}$. Note that (IC6) holds in the binary case, though.

We do not discuss the majority or arbitration postulates here (except to note that majority is easily handled by a straightforward modification to Definition 5); this is discussed in the full paper. Note however that the present approach satisfies a non-majority postulate, viz.:

$$\Delta^\mu(\Psi_1 \sqcup \Psi_2^n) = \Delta^\mu(\Psi_1 \sqcup \Psi_2).$$

This postulate is identified in [11], a weaker version of which is used to define their arbitration operator.

Theorem 2. Let Δ and Δ_c be defined as in Definition 5.

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)' $\alpha \Delta \beta$ is satisfiable iff α is satisfiable and β is satisfiable.
 (LS6)' $(\alpha \Delta \beta_1) \wedge \beta_2$ implies $\alpha \Delta (\beta_1 \wedge \beta_2)$.
 (LS8)' If α is satisfiable and β is satisfiable then $\alpha \wedge (\alpha \Delta \beta)$ is satisfiable.
3. (LS6c)' For any selection function c there is a selection function c' such that $\alpha \Delta_c \beta_1$ implies $\alpha \Delta_{c'} (\beta_1 \vee \beta_2)$ or $\alpha \Delta_c \beta_2$ implies $\alpha \Delta_{c'} (\beta_1 \vee \beta_2)$.

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

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

We get that:

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

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

Theorem 3. Let $\dot{+}$ and Δ be given as in Definitions 2 and 5 (respectively). Then,

1. $\alpha \Delta \beta = \alpha \dot{+} \beta \cap \beta \dot{+} \alpha$.
2. $\alpha \dot{+} \beta = \alpha \Delta^\beta \top$.

⁶ 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.

3.3 Belief Set Projection

In our second approach, the contents of several belief sets are “projected” onto another.⁷ Again, the formulation is straightforward within the framework of belief change scenarios. For belief sets K_1, \dots, 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 again represent a set of entailment-based and consistency-based integrity constraints, respectively.

Definition 6. 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

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

Then

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

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” the belief set onto the formula that we revise with. Similarly, the actual projection operation “projects” the belief sets onto whatever is contained in R .

Definition 7 (Merging via Projection). Let \mathcal{K} be a family of sets of formulas in $\mathcal{L}_{\mathcal{P}}$ and R and C be finite 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 4. We have that $(p \wedge q \wedge r) \nabla (p \wedge \neg q)$ yields two EQ sets:

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

$$EQ_2 = \{p^1 \equiv p, p^2 \equiv p, q^2 \equiv q, r^1 \equiv r, r^2 \equiv r\}.$$

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

⁷ We thank Jérôme Lang for pointing out this alternative to us.

Example 5. Consider the example from [11]:

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

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

$$K_1 \nabla K_2 = Ch(\top).$$

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 4. *Let ∇ and ∇_c be defined as in Definition 7.*

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

(IC1') *If $\bigwedge \Psi \not\vdash \neg\mu$ and $\mu \not\vdash \perp$ then $\nabla^\mu(\Psi) \not\vdash \perp$.*⁸

(IC4') *If $\phi_1 \not\vdash \perp$, $\phi_2 \not\vdash \perp$ and $\phi_1 \vdash \mu$ and $\phi_2 \vdash \mu$ then: $\nabla^\mu(\phi_1 \sqcup \phi_2) \wedge \phi_1 \not\vdash \perp$.*

Theorem 5. *Let ∇ and ∇_c be defined as in Definition 7.*

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):

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

does not imply

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

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 project in our framework.

Further, postulate (LS7) does not hold here, as Example 5 illustrates, nor is the projection operator associative.

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

Theorem 6. *Let \mathcal{K} , $\Delta^{R,C}$ and $\nabla^{R,C}$ be given as in Definitions 5 and 7 (respectively).*

$$\nabla^{R,C}(\mathcal{K}) \subseteq \Delta^{R,C}(\mathcal{K}).$$

That is, in binary terms, $\alpha \nabla^{R,C} \beta \subseteq \alpha \Delta^{R,C} \beta$.

As well, we have the following analogue to Theorem 3:

Theorem 7. *Let $\dot{+}$ and ∇ be given as in Definitions 2 and 7 (respectively).*

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

4 Complexity

In [6], we analysed the computational complexity of reasoning from belief change scenarios. Specifically, we addressed the following basic reasoning tasks:

Theorem 8 ([6]).

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

5 Discussion

We have presented two approaches for merging belief sets, expressed in a general, consistency-based framework for belief change [5]. 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 at the level of belief sets rather than models. Thus projection can be regarded as an operator lying intermediate between belief revision and update.

In the full paper we consider merging and projection with respect to a denumerable number of belief sets. As well, we show how these operations (in the finite case) can be equivalently expressed as functions with domain and range effectively being knowledge bases, that is, arbitrary subsets of \mathcal{L} , while retaining syntax-independence. Last, we provide abstract algorithms for computing these operators.

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 for their detailed and helpful comments.

References

1. C. Baral, S. Kraus, J. Minker, and V. Subrahmanian. Combining multiple knowledge bases consisting of first order theories. *Computational Intelligence*, 8(1):45–71, 1992.
2. S. Benferhat, D. Dubois, S. Kaci, and H. Prade. Possibilistic merging and distance-based fusion of propositional information. *Annals of Mathematics and Artificial Intelligence*, 34(1-3):217–252, 2003.
3. R. Booth. Social contraction and belief negotiation. In D. Fensel, F. Giunchiglia, D. McGuinness, and M. Williams, editors, *Proceedings of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning*, pages 375–384, San Francisco, 2002. Morgan Kaufmann.
4. M. Dalal. Investigations into theory of knowledge base revision. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 449–479, St. Paul, Minnesota, 1988.
5. J. Delgrande and T. Schaub. A consistency-based approach for belief change. *Artificial Intelligence*, 151(1-2):1–41, 2003.
6. J.P. Delgrande, T. Schaub, H. Tompits, and S. Woltran. On computing solutions to belief change scenarios. In S. Benferhat and P. Besnard, editors, *Proceedings of the Sixth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2001)*, volume 2143 of *Lecture Notes in Artificial Intelligence*, pages 510–521, Toulouse, Fr., 2001. Springer Verlag.
7. P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.
8. S. Konieczny, J. Lang, and P. Marquis. Distance-based merging: a general framework and some complexity results. In D. Fensel, F. Giunchiglia, D. McGuinness, and M. Williams, editors, *Proceedings of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning*, pages 97–108, San Francisco, 2002. Morgan Kaufmann.
9. S. Konieczny and R. Pino Pérez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

10. Sébastien Konieczny. On the difference between merging knowledge bases and combining them. In A. G. Cohn, F. Giunchiglia, and B. Selman, editors, *KR2000: Principles of Knowledge Representation and Reasoning*, pages 135–144, San Francisco, 2000. Morgan Kaufmann.
11. Sébastien Konieczny and Ramón Pino Pérez. On the logic of merging. In A. G. Cohn, L. Schubert, and S. C. Shapiro, editors, *KR'98: Principles of Knowledge Representation and Reasoning*, pages 488–498. Morgan Kaufmann, San Francisco, California, 1998.
12. P. Liberatore and M. Schaerf. Reducing belief revision to circumscription (and vice versa). *Artificial Intelligence*, 93(1–2):261–296, 1997.
13. P. Liberatore and M. Schaerf. Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering*, 10(1):76–90, 1998.
14. J. Lin and A. Mendelzon. Knowledge base merging by majority. In *Dynamic Worlds: From the Frame Problem to Knowledge Management*. Kluwer, 1999.
15. T. Meyer. On the semantics of combination operations. *Journal of Applied NonClassical Logics*, 11(1-2):59–84, 2001.
16. R. Reiter. Towards a logical reconstruction of relational database theory. In M.L. Brodie, J. Mylopoulos, and J.W. Schmidt, editors, *On Conceptual Modelling*, pages 191–233. Springer-Verlag, 1984.
17. P. Revesz. On the semantics of theory change: Arbitration between old and new information. In C. Beeri, editor, *Proceedings of the Twelfth ACM Symposium on Principles of Database Systems*, pages 71–82, Washington D.C., 1993.
18. F. Sadri and R. Kowalski. A theorem-proving approach to database integrity. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*, chapter 9, pages 313–362. Morgan Kaufmann Publishers, 1987.
19. W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume II, pages 105–134. Kluwer Academic Publishers, 1988.