

Consistency-Based Approaches to Merging Knowledge Bases: Preliminary Report

James P. Delgrande

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

Torsten Schaub*

Institut für Informatik,
Universität Potsdam
D-14482 Potsdam,
Germany
torsten@cs.uni-potsdam.de

Abstract

We present a framework for investigating 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 K_1 and K_2 is a maximal consistent (if possible) set of formulas comprising the joint knowledge of K_1 and K_2 . 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: Two knowledge sources are "projected" onto a third (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, interrelationships between these approaches and belief revision are given.

Introduction

The problem of merging multiple, potentially conflicting, sources or bodies of information arises in various guises. For example, an intelligent agent may receive reports from differing sources of knowledge that must be combined. For example, 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 the 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 (fol-

lowing (Konieczny & Pino Pérez 1998)) *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. (Konieczny & Pino Pérez 1998) characterises these approaches as trying to minimize global dissatisfaction vs. trying to minimize local dissatisfaction, respectively.

In this paper, we develop a specific framework for specifying merge operations. This framework extends our earlier work in belief revision. There, and here, 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 two 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. In the second approach, knowledge sources are projected onto a third knowledge source (which in the simplest case would consist solely of \top). That is, the two sources we wish to merge are used to augment the knowledge of a third source. This second approach then appears to be a natural extension to belief *revision*. In both cases, 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.

Due to space limitations, we present the approaches in their most basic form. While the approaches extend naturally and straightforwardly to any (denumerable) number of knowledge bases, here we simply deal with binary merge operators. As well, while the approaches admit a simple abstract implementation, we defer the development of specific algorithms to a later paper. The next section describes related work while Section develops our approaches. We conclude with a discussion.

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

Background

Belief Revision

A standard approach in belief revision is to provide a set of *rationality postulates* for a belief revision function. The *AGM approach* of Alchourrón, Gärdenfors, and Makinson (Alchourrón, Gärdenfors, & Makinson 1985; Gärdenfors 1988) provides the best-known set of such postulates. The approach assumes a language \mathcal{L} , closed under the usual set of Boolean connectives; the language is assumed to be governed by a logic that includes classical propositional logic, and that is compact. Belief change is described at the *knowledge level*, that is on an abstract level, independent of how beliefs are represented and manipulated. Belief states are modelled by logically closed sets of sentences, called *belief sets*. So a belief set K can be seen as a partial theory of the world. For belief set K and formula α , $K + \alpha$ is the deductive closure of $K \cup \{\alpha\}$, called the *expansion* of K by α . K_{\perp} is the inconsistent belief set (i.e. K_{\perp} is the set of all formulas). A *revision* function $\dot{+}$ is a function from $2^{\mathcal{L}} \times \mathcal{L}$ to $2^{\mathcal{L}}$. The AGM approach proposes eight postulates that revision function $\dot{+}$ should satisfy; for details we refer to (Alchourrón, Gärdenfors, & Makinson 1985).

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_{\perp} , \dots ¹. Thus $K = \text{Cn}(K)$, where $\text{Cn}(\cdot)$ 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}'). This is defined analogously for sets of formulas.

A *belief change scenario* in $\mathcal{L}_{\mathcal{P}}$ was defined in (Delgrande & Schaub 2003) as a triple $B = (K, R, C)$ where K , R , and C are sets of formulas in $\mathcal{L}_{\mathcal{P}}$. Informally, K is a knowledge base 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

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

Then

$$\text{Cn}(K' \cup R \cup EQ) \cap \mathcal{L}_{\mathcal{P}}$$

¹We note in passing though that all our definitions, following, work for arbitrary sets of formulas, and provide the basis for a finite representation of these operators.

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. 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 . When we come to define revision in Definition 2, we will use a selection function to select a specific consistent belief change extension. This use of selection functions then is slightly different from that in the AGM approach.

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. In the definitions below, note that K need not be a belief set, but rather may be any arbitrary set of formulas.

Definition 2 (Revision) Let K be a knowledge base 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 .

Table 1 gives examples of skeptical revision. The first column specifies the original knowledge base, but with atoms already renamed. The second column gives the revision formula, while the third lists the determining EQ set(s), and the last column gives the results of the revision. For the first and last column, we give a formula whose deductive closure is the corresponding belief set.

K'	α	EQ	$K \dot{+} \alpha$
$p' \wedge q'$	$\neg q$	$\{p \equiv p'\}$	$p \wedge \neg q$
$\neg p' \equiv q'$	$\neg q$	$\{p \equiv p', q \equiv q'\}$	$p \wedge \neg q$
$p' \vee q'$	$\neg p \vee \neg q$	$\{p \equiv p', q \equiv q'\}$	$p \equiv \neg q$
$p' \wedge q'$	$\neg p \vee \neg q$	$\{p \equiv p'\}, \{q \equiv q'\}$	$p \equiv \neg q$

Table 1: Skeptical revision examples.

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

Definition 1 also leads to a natural and general treatment of integrity constraints. There are two standard definitions of a knowledge base K satisfying a static integrity constraint IC . In the *consistency-based* approach of (Kowalski 1978;

Sadri & Kowalski 1987), K satisfies IC iff $K \cup \{IC\}$ is satisfiable. In the *entailment-based* approach of (Reiter 1984), K satisfies IC iff $K \vdash IC$. In our approach, we define revision taking into account both approaches to integrity constraints. Corresponding to Definition 2 (and ignoring the choice approach) we obtain:

Definition 3 Let K be a knowledge base, α a formula, and IC_e, IC_c sets of formulas. Let $(E_i)_{i \in I}$ be the family of all belief change extensions of $(K, \{\alpha\} \cup IC_e, \overline{IC_c})$ where $\overline{IC_c} = \{\neg\delta \mid \delta \in IC_c\}$.

Then, we define $K \dot{+}^{(IC_e, IC_c)} \alpha = \bigcap_{i \in I} E_i$ as the revision of K by α incorporating integrity constraints IC_e (entailment-based) and IC_c (consistency-based).

It proves to be the case that following a revision, the entailment-based constraints are true in the resulting belief set, and the consistency-based constraints are consistent.

Belief Merging

Konieczny and Pino Perez (Konieczny & Pino Pérez 2002) consider the problem of merging knowledge bases coming from different sources. To this end, they consider multisets of the form $\Psi = \{K_1, \dots, K_n\}$ and assume that all knowledge bases K_i are consistent, finite, 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 (Konieczny & Pino Pérez 2002),² we use $\Delta^\mu(\Psi)$ 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 4 ((Konieczny & Pino Pérez 2002)) Let Ψ be a multiset, 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 $\Delta^\mu(\Psi) \not\vdash \perp$.³
- (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

²(Konieczny & Pino Pérez 1998) write $\Delta_\mu(\Psi)$ where we have $\Delta^\mu(\Psi)$.

³ Ψ is consistent just if $\Psi = \{K_1, \dots, K_n\}$ and $K_1 \cup \dots \cup K_n \not\vdash \perp$.

integrity constraints, when consistent. (IC4) is a *fairness* postulate, that when two knowledge 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 (Konieczny & Pino Pérez 1998).

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 Ψ_2 , this multiset will eventually come to dominate the merge operation.

An arbitration operator is characterised by the original postulates together with the postulate:

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

Examples are given of a merging operator using Dalal's notion of distance (Dalal 1988).

Liberatore and Schaerf consider merging two knowledge bases in (Liberatore & Schaerf 1998) and propose the following postulate set to characterise a merge operator that they call an *arbitration* operator and (Konieczny & Pino Pérez 2002) calls a *commutative revision operator*. Like (Konieczny & Pino Pérez 2002) 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 (Baral *et al.* 1992) and (Revesz 1993). The former proposes various theory merging operators based on the selection of maximum consistent subsets in the union of the knowledge bases; see (Konieczny 2000) for a pertinent discussion. The latter proposes an "arbitration" operator that satisfies a subset of the Liberatore and Schaerf postulates; see (Liberatore & Schaerf 1997) for a discussion. (Lin & Mendelzon 1999) first identified and addressed the majority merge operator. (Konieczny, Lang, & Marquis 2002) 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, (Meyer 2001) proposed a general approach to formulating merging functions, based on ordinal conditional functions (Spohn 1988). 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. (Booth 2002) 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 knowledge base.

Consistency-Based Approaches to Belief Set Merging

In this section we modify the framework given by Definition 1 to deal with knowledge base *merging*, in which multiple sources of information (knowledge bases, etc.) are coalesced into a single knowledge base. While we just consider merging two knowledge sets, the more general problem of merging an arbitrary number of knowledge sets is a straightforward extension. We detail two distinct approaches to knowledge base merging, expressible in the general approach.

In the first case, the intuition is that for merging two knowledge bases, the common information is in a sense “pooled”. This approach then seems to conform to the commonsense notion of merging of knowledge, in which two sets of knowledge are joined to produce a single knowledge base retaining as much as possible of the contents of the original knowledge sets. In the second approach, two knowledge sources are projected onto a third knowledge source (which in the simplest case could consist solely of \top). That is, the two sources we wish to merge are used to augment the knowledge of a third source. This (general) approach appears to differ from others that have appeared in the literature; as well, as we show, it appears to us that it has reasonable properties with respect to commonsense notions of merging.

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 knowledge bases 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 , except that the single set of formulas K is extended to several of sets of formulas.

In this paper, we consider only the case of two knowledge bases, that is, where $|J| = 2$. We denote such binary be-

lief change scenarios as $(\{K_1, K_2\}, R, C)$. Extending the approaches to the more general cases is carried out in a subsequent paper. 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 in the result. Hence, the negations of the formulas in C will (individually) be consistent with the result of a merge operation. While R is intended to represent a set of entailment-based integrity constraints (Reiter 1984), 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 (Kowalski 1978; Sadri & Kowalski 1987), 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 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 . As well, 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 sections.

Belief Set Merging

Consider the first approach, in which the contents of belief sets are to be merged. Since we are assuming binary belief change scenarios, we will write the merge operator Δ as an infix operator.

Definition 5 Let $B = (\{K_1, K_2\}, R, C)$ be a multi belief change scenario in $\mathcal{L}_{\mathcal{P}}$. Define EQ as a maximal set of equivalences $EQ \subseteq \{p^1 \equiv p^2 \mid p \in \mathcal{P}\}$ such that^{4, 5}

$$Cn(K_1^1 \cup K_2^2 \cup R^{1,2} \cup EQ) \cap (C^{1,2} \cup \{\perp\}) = \emptyset$$

Then

$$\{\alpha \mid \alpha^1, \alpha^2 \in Cn(K_1^1 \cup K_2^2 \cup R^{1,2} \cup EQ)\}$$

is a symmetric 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) symmetric belief change extension of B .

We use $R^{1,2}$ rather than R in the above definition, since the integrity constraints must be true in both (relabelled) belief sets K_1^1 and K_2^2 ; a similar comment applies to $C^{1,2}$.

For simplicity, we give the next definition in terms of formulas, bearing in mind that in our setting a finite knowledge base can be identified with the conjunction of its elements.

⁴Another way of expressing this condition is that $K_1^1 \cup K_2^2 \cup R^{1,2} \cup EQ \not\vdash \phi$ for every $\phi \in C^{1,2} \cup \{\perp\}$.

⁵A more explicit way, stressing that the superscript denotes an operation, would be to write $(K_1^1)^1$ instead of K_1^1 . For example, given $K_1 = \{p \wedge q\}$, we get $K_1^1 = \{p^1 \wedge q^1\}$.

Definition 6 (Merging) Let α and β be formulas and let $(E_i)_{i \in I}$ be the family of all symmetric belief change extensions of $(\{\alpha\}, \{\beta\})$, R, C .

Then, we define

1. $\alpha \Delta_c^{R,C} \beta = E_i$ as the choice merging of α and β with respect to some selection function c with $c(I) = i$.
2. $\alpha \Delta^{R,C} \beta = \bigcap_{i \in I} E_i$ as the (skeptical) merging of α and β .

In the above, if $R = C = \emptyset$, then we just write $\alpha \Delta \beta$. For complying with the notation used in Definition 4 (Konieczny & Pino Pérez 2002), 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 $Ch(\{p \wedge r \wedge s\})$.

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 $EQ = \{\}$ and in fact

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

This example is introduced and discussed in (Konieczny & Pino Pérez 1998); 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 (Konieczny & Pino Pérez 1998) 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.⁶ 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 .

Theorem 1 Let Δ^μ and Δ_c^μ be defined as in Definition 6 and let $|\Psi| = 2$.

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

(IC1') If $\Psi = \{K_1, K_2\}$, $K_1 \wedge \mu \not\vdash \perp$, $K_2 \wedge \mu \not\vdash \perp$ then $\Delta^\mu(\Psi) \not\vdash \perp$.⁷

⁶(Konieczny & Pino Pérez 1998) go on to suggest that the appropriate belief is that precisely that two stocks will rise, and two not. We do not agree with this analysis, but rather it seems that the most appropriate merged belief set will allow for each stock that it will rise or not (i.e. nothing of substance is believed).

⁷It is straightforward to obtain (IC1) by essentially ignoring belief sets that are inconsistent with μ . We remain with the present postulate since it reflects the most natural formulation of merge in our framework.

Recall that we are addressing the case of two belief sets only. As well, in treating Δ as an infix operator rather than an operator on sets, we have a commutativity postulate also satisfied in our approach, that is satisfied trivially in (Konieczny & Pino Pérez 2002), viz.:

$$K_1 \Delta^\mu K_2 = K_2 \Delta^\mu K_1.$$

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); again, see the full paper.

Theorem 2 Let Δ and Δ_c be defined as in Definition 6.

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.

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

We get that:

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

Lastly, we have the following result showing that in this approach, merging two knowledge bases 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 6 (respectively). Then,

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

Belief Set Projection

In our second approach, the contents of two belief sets are “projected” onto a third. Again, the formulation is straightforward within the framework of belief change scenarios. For belief sets K_1 and K_2 we express each belief set in a distinct language, but project these two belief sets onto a third belief set in which R is believed. (In the simplest case we would have $R \equiv \top$.)

Definition 7 Let $B = (\{K_1, K_2\}, R, C)$ be a multi belief change scenario in $\mathcal{L}_{\mathcal{P}}$. Define EQ as a maximal set of equivalences

$$EQ \subseteq \{p^i \equiv p \mid p \in \mathcal{P} \text{ and } i \in \{1, 2\}\}$$

such that⁸

$$Ch(K_1^1 \cup K_2^2 \cup R \cup EQ) \cap (C \cup \{\perp\}) = \emptyset$$

Then

$$Ch(K_1^1 \cup K_2^2 \cup R \cup EQ) \cap \mathcal{L}_{\mathcal{P}}$$

is a projected 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) symmetric belief change extension of B .

There is an interesting similarity between revision and projection. Revision in some sense “projects” the knowledge base onto the formula that we revise with. Similarly, the actual projection operation “projects” two knowledge bases onto whatever is contained in R .

Definition 8 (Merging via Projection) Let α and β be formulas and let $(E_i)_{i \in I}$ be the family of all projected belief change extensions of $(\{\alpha\}, \{\beta\}, R, C)$.

Then, we define

1. $\alpha \nabla_c^{R,C} \beta = E_i$ as the projection of α and β with respect to selection function c with $c(I) = i$.
2. $\alpha \nabla^{R,C} \beta = \bigcap_{i \in I} E_i$ as the (skeptical) projection of α and β onto R .

As above, 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:

$$\begin{aligned} EQ_1 &= \{p^1 \equiv p, p^2 \equiv p, q^1 \equiv q, r^1 \equiv r, r^2 \equiv r\} \\ &\text{and} \\ EQ_2 &= \{p^1 \equiv p, p^2 \equiv p, q^2 \equiv q, r^1 \equiv r, r^2 \equiv r\}. \end{aligned}$$

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

Example 5 Consider the example from (Konieczny & Pino Pérez 1998):

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

In forming 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 obtain that

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

We feel that this is the correct 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

⁸Or: $K_1^1 \cup K_2^2 \cup R \cup EQ \not\vdash \phi$ for every $\phi \in C \cup \{\perp\}$.

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 would 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 and let $|\Psi| = 2$.

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

(IC1’) If $\bigwedge \Psi \not\vdash \neg \mu$ and $\mu \not\vdash \perp$ then $\nabla^\mu(\Psi) \not\vdash \perp$.⁹

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

Theorem 5 Let ∇ and ∇_c be defined as in Definition 8.

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

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

Further, postulate (LS7) does not hold here, as Example 5 illustrates.

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

Theorem 6 Let $\Delta^{R,C}$ and $\nabla^{R,C}$ be given as in Definitions 6 and 8 (respectively).

$$\alpha \nabla^{R,C} \beta \subseteq \alpha \Delta^{R,C} \beta.$$

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

Theorem 7 Let $\dot{+}$ and Δ be given as in Definitions 2 and 8 (respectively).

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

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

Complexity

In (Delgrande *et al.* 2001), we analysed the computational complexity of reasoning from belief change scenarios. Specifically, we addressed with the following basic reasoning tasks:

DEFEXT: Decide whether a belief change scenario B has a consistent belief change extension.

CHOICE: Given a belief change scenario B and formula ϕ , decide whether ϕ is contained in at least one consistent belief change extension of B .

SKEPTICAL: Given a belief change scenario B and formula ϕ , decide whether ϕ is contained in all consistent belief change extensions of B .

We obtained the following consistency results.

Theorem 8 (Delgrande *et al.* 2001) *We have the following completeness results:*

1. DEFEXT is NP -complete;
2. CHOICE is Σ_P^2 -complete; and
3. SKEPTICAL is Π_P^2 -complete.

Clearly, the variants of these decision problems for binary 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.

Discussion

We have presented two approaches for merging knowledge bases, expressed in a general, consistency-based framework for belief change (Delgrande & Schaub 2003). In the first approach, the intuition is that for merging two knowledge bases, common information is in a sense “pooled”. This approach then seems to conform to the commonsense notion of merging of knowledge, in which two sets of knowledge are joined to produce a single knowledge base retaining as much as possible of the contents of the original knowledge sets. In the second approach, two knowledge sources are projected onto a third knowledge source. That is, the two sources we wish to merge are used to augment the knowledge of a third (possibly empty) source. 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 difficulty illustrated in Example 2.

Additionally, the second approach, projection, has something of the flavour of both belief revision and update. With respect to belief revision, projection can be viewed as a process whereby two knowledge bases are simultaneously revised (and consequently merged) with respect to a third. With respect to belief update, in an update operation, individual models of a belief set are updated by the sentence for update. 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.

The general approach can be straightforwardly extended in several directions. In the full paper we consider merging and projection with respect to a denumerable number of knowledge bases. 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.

Acknowledgments 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 thank the two referees for their detailed and helpful comments.

References

- Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50(2):510–530.
- Baral, C.; Kraus, S.; Minker, J.; and Subrahmanian, V. 1992. Combining multiple knowledge bases consisting of first order theories. *Computational Intelligence* 8(1):45–71.
- Booth, R. 2002. Social contraction and belief negotiation. In Fensel, D.; Giunchiglia, F.; McGuinness, D.; and Williams, M., eds., *Proceedings of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning*, 375–384. San Francisco: Morgan Kaufmann.
- Dalal, M. 1988. Investigations into theory of knowledge base revision. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, 449–479.
- Delgrande, J., and Schaub, T. 2003. A consistency-based approach for belief change. *Artificial Intelligence* 151(1-2):1–41.
- Delgrande, J.; Schaub, T.; Tompits, H.; and Woltran, S. 2001. On computing solutions to belief change scenarios. In Benferhat, S., and Besnard, P., eds., *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*, 510–521. Toulouse, Fr.: Springer Verlag.
- Gärdenfors, P. 1988. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. Cambridge, MA: The MIT Press.
- Konieczny, S., and Pino Pérez, R. 1998. On the logic of merging. In Cohn, A. G.; Schubert, L.; and Shapiro, S. C., eds., *KR'98: Principles of Knowledge Representation and Reasoning*. San Francisco, California: Morgan Kaufmann. 488–498.
- Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: A logical framework. *Journal of Logic and Computation* 12(5):773–808.

Konieczny, S.; Lang, J.; and Marquis, P. 2002. Distance-based merging: a general framework and some complexity results. In Fensel, D.; Giunchiglia, F.; McGuinness, D.; and Williams, M., eds., *Proceedings of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning*, 97–108. San Francisco: Morgan Kaufmann.

Konieczny, S. 2000. On the difference between merging knowledge bases and combining them. In Cohn, A. G.; Giunchiglia, F.; and Selman, B., eds., *KR2000: Principles of Knowledge Representation and Reasoning*, 135–144. San Francisco: Morgan Kaufmann.

Kowalski, R. 1978. Logic for data description. In Gallaire, H., and Minker, J., eds., *Logic and Data Bases*. Plenum Press. 77–103.

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

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

Lin, J., and Mendelzon, A. 1999. Knowledge base merging by majority. In *Dynamic Worlds: From the Frame Problem to Knowledge Management*. Kluwer.

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

Reiter, R. 1984. Towards a logical reconstruction of relational database theory. In Brodie, M.; Mylopoulos, J.; and Schmidt, J., eds., *On Conceptual Modelling*. Springer-Verlag. 191–233.

Revesz, R. 1993. On the semantics of theory change: Arbitration between old and new information. In Beeri, C., ed., *Proceedings of the Twelfth ACM Symposium on Principles of Database Systems*, 71–82.

Sadri, F., and Kowalski, R. 1987. A theorem-proving approach to database integrity. In Minker, J., ed., *Foundations of Deductive Databases and Logic Programming*. Morgan Kaufmann Publishers. chapter 9, 313–362.

Spohn, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, volume II. Kluwer Academic Publishers. 105–134.