A Consistency-Based Approach for Belief Change

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

This paper presents a general, consistency-based framework for expressing belief change. The framework has good formal properties while being well-suited for implementation. For belief revision, informally, in revising a knowledge base K by a sentence α, we begin with α and include as much of K as consistently possible. This is done by expressing K and α in disjoint languages, asserting that the languages agree on the truth values of corresponding atoms wherever consistently possible, and then re-expressing the result in the original language of K. There may be more than one way in which the languages of K and α can be so correlated: in *choice revision*, one such "extension" represents the revised state; alternately (skeptical) *revision* consists of the intersection of all such extensions. Contraction is similarly defined although, interestingly, it is not interdefinable with revision.

The framework is general and flexible. For example, one could go on and express other belief change operations such as update and erasure, and the the merging of knowledge bases. Further, the framework allows the incorporation of static and dynamic integrity constraints. The approach is well-suited for implementation: belief change can be equivalently expressed in terms of a finite knowledge base; and the scope of a belief change operation can be restricted to just those propositions common to the knowledge base and sentence for change. We give a high-level algorithm implementing the procedure, and an expression of the approach in Default Logic. Lastly, we briefly discuss two implementations of the approach.

Keywords: Belief change, belief revision and contraction, consistency-based reasoning

1 Introduction

This paper describe a general framework for expressing belief change, focussing on revision and contraction. A key feature of the framework is that it combines theoretical and practical considerations in a single system: revision and contraction operators have good formal properties (satisfying most AGM postulates) while being well-suited for implementation. Informally, to revise a knowledge base K by sentence α, we begin with α and "include" as much of K as consistently possible.

This is carried out by expressing K and α in disjoint languages, "forcing" (via a maximisation process) the languages to agree on truth values of atoms wherever consistently possible, and then re-expressing the result in the original language of K. There may be more than one way in which the maximisation process can be carried out. This inherent non-determinism gives rise to two notions of revision. In "choice" revision one such "extension" is selected for the revised state. In general "skeptical" revision, the revised state consists of the intersection of all such extensions. Belief contraction is defined analogously. Since we are maximising equivalences over a set of atomic sentences, the approach has the same flavour as the consistency-based approaches for diagnosis [Rei87b], or default reasoning [Poo88], or assumption-based truth maintenance [RdK87].

The approach is developed first in a formal, abstract framework. The central notion is that of a *belief change scenario* consisting of a triple of sets of formulas, $B = (K, R, C)$. Informally, K is a knowledge base that will be changed such that the set R will be derivable in the resulting knowledge base, while members of C will not. Revision and contraction are then easily defined, by letting $C = \emptyset$ and $R = \emptyset$, respectively. Update, erasure, and merging are similarly definable although we do not do so here. Moreover it is straightforward to incorporate different sorts of integrity constraints in this framework.

The approach is independent of syntax, in that revising (or contracting) a knowledge base K by sentence α is independent of how K and α are expressed. The belief change operators are also shown to satisfy the majority of the AGM postulates, with the exception of a "non-basic" postulate and, in the case of contraction, the recovery postulate. On the other hand, the approach is well-suited for implementation. Belief change can be expressed in terms of a finite knowledge base, in place of a deductively-closed belief set. Further, the scope of a belief change operator can be restricted to those propositions common to a knowledge base and sentence for change. We provide a high-level algorithm implementing the approach, and show how the approach can be expressed using Default Logic [Rei80]. Finally we briefly describe two implementations of the approach.

In the next section we briefly review approaches to belief change. In Section 3 we discuss intuitions underlying our approach and, in particular, the suitability of a consistency-based approach. Section 4 presents the general framework, then explores revision and contraction. In Section 5 we consider implementation issues, while in Section 6 we compare our approach with related work. We conclude in Section 7 with a summation and discussion. Proofs of theorems are contained in an appendix. In [DS02], we further explore the general framework, and show that it is flexible enough to express other belief change operations such as update, erasure, and merging.

2 Background

A common approach in addressing belief change has been to provide a set of *rationality postulates* for a belief change function. These rationality postulates constrain, or give properties of, such functions, but have little to say about how a specific function is to be implemented. The *AGM approach* of Alchourron, Gärdenfors, and Makinson [AGM85, Gär88] provides the best-known set of such postulates; see also [Han99, Rot01] for extensive discussions of this and other approaches. 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*. Thus, a belief set is a set K of sentences which satisfies the constraint:

If K logically entails β then $\beta \in K$.

So 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 $2^{\mathcal{L}}$ satisfying the following postulates.

$(K\dot{+}1)$ $K\dot{+}\alpha$ is a belief set.

$(K\dot{+}2)$ $\alpha \in K\dot{+}\alpha$.

$(K\dot{+}3)$ $K\dot{+}\alpha \subseteq K + \alpha$.

$(K\dot{+}4)$ If $\neg\alpha \notin K$, then $K + \alpha \subseteq K\dot{+}\alpha$.

$(K\dot{+}5)$ $K\dot{+}\alpha = K_\perp$ iff $\vdash \neg\alpha$.

$(K\dot{+}6)$ If $\vdash \alpha \equiv \beta$, then $K\dot{+}\alpha = K\dot{+}\beta$.

$(K\dot{+}7)$ $K\dot{+}(\alpha \wedge \beta) \subseteq (K\dot{+}\alpha) + \beta$.

$(K\dot{+}8)$ If $\neg\beta \notin K\dot{+}\alpha$, then $(K\dot{+}\alpha) + \beta \subseteq K\dot{+}(\alpha \wedge \beta)$.

That is: the result of revising K by α is a belief set in which α is believed; whenever the result is consistent, revision consists of the expansion of K by α; the only time that $K_\perp$ is obtained is when $\neg\alpha$ is a tautology; and revision is independent of the syntactic form of K and α. The last two postulates deal with the relation between revising with a conjunction and expansion.

Contraction is the dual notion of revision, in which beliefs are retracted but no new beliefs are added. In the AGM approach, a contraction function $\dot{-}$ is a function from $2^{\mathcal{L}} \times \mathcal{L}$ to $2^{\mathcal{L}}$ satisfying the following postulates.

$(K\dot{-}1)$ $K\dot{-}\alpha$ is a belief set.

$(K\dot{-}2)$ $K\dot{-}\alpha \subseteq K$.

$(K\dot{-}3)$ If $\alpha \notin K$, then $K\dot{-}\alpha = K$.

$(K\dot{-}4)$ If $\not\vdash \alpha$, then $\alpha \notin K\dot{-}\alpha$.

$(K\dot{-}5)$ If $\alpha \in K$, then $K \subseteq (K\dot{-}\alpha) + \alpha$.

$(K\dot{-}6)$ If $\vdash \alpha \equiv \beta$, then $K\dot{-}\alpha = K\dot{-}\beta$.

$(K\dot{-}7)$ $K\dot{-}\alpha \cap K\dot{-}\beta \subseteq K\dot{-}(\alpha \wedge \beta)$.

$(K\dot{-}8)$ If $\beta \notin K\dot{-}(\alpha \wedge \beta)$, then $K\dot{-}(\alpha \wedge \beta) \subseteq K\dot{-}\beta$.

Revision and contraction are often interdefinable by means of the following identities:

Levi Identity: $K\dot{+}\alpha = (K\dot{-}\neg\alpha) + \alpha$.

Harper Identity: $K\dot{-}\alpha = K \cap (K\dot{+}\neg\alpha)$.

The Levi Identity asserts that revision by α corresponds to contraction by $\neg\alpha$ followed by expansion by α, while the Harper Identity asserts that contracting K by α corresponds to selecting just those sentences of K that remain if K is revised by $\neg\alpha$.

Various constructions based on *preference relations* have been proposed, in terms of which belief change functions can be defined. Earliest and best-known among these is *epistemic entrenchment orderings* [Gär88]. An epistemic entrenchment ordering related to a belief set K is a binary relation $\leq$ on the formulas in $\mathcal{L}$, reflecting the relative degree of acceptance of sentences. Belief change can also be characterised by a total preorder on interpretations in the language [Gro88].

The postulate sets for belief change, and their accompanying constructions, do not address the issue of *iterated* belief revision. However, clearly, one would be interested in not just a single revision of a belief set by a formula, but also in sequences of revisions. [Leh95] provides an extended set of rationality postulates; other representative work includes [BG93, Bou94, Wil94, NFPS96, DP97, Pap01]. Much, if not all, of this work is based upon or inspired by [Spo88]. However, it has proven to be very difficult to develop a belief revision operator with plausible properties for iterated revision; see [NFPS96, DP97] for excellent discussions. We briefly discuss Darwiche and Pearl's approach here, as a more recent and well-known proposal.

Darwiche and Pearl employ the notion of an *epistemic state* that encodes how a revision function changes following a revision. They propose the following postulates.[1]

$(C1)$ If $\alpha \vdash \beta$ then $(K\dot{+}\beta)\dot{+}\alpha = K\dot{+}\alpha$.

$(C2)$ If $\alpha \vdash \neg\beta$ then $(K\dot{+}\beta)\dot{+}\alpha = K\dot{+}\alpha$.

$(C3)$ If $\beta \in K\dot{+}\alpha$ then $\beta \in (K\dot{+}\beta)\dot{+}\alpha$.

$(C4)$ If $\neg\beta \notin K\dot{+}\alpha$ then $\neg\beta \notin (K\dot{+}\beta)\dot{+}\alpha$.

[NFPS96] propose a variant of $(C2)$ along with the following postulate:

$(Conj)$ If $\alpha \wedge \beta \nvdash \bot$ then $(K\dot{+}\alpha)\dot{+}^{\alpha}\beta = K\dot{+}(\alpha \wedge \beta)$.

The superscript on $\dot{+}^{\alpha}$ indicates that following revision by α, $\dot{+}$ depends in part on α. This postulate is strong enough to derive $(C1)$, $(C3)$, and $(C4)$ in the presence of the AGM postulates.

These postulates are not uncontentious. For example, an instance of $(C2)$ (letting α be $\neg p$ and β be $p \wedge q$) is the following:

$(C2')$ $(K\dot{+}(p \wedge q))\dot{+}\neg p \quad \equiv \quad K\dot{+}\neg p.$

[1] Darwiche and Pearl phrase their postulates in terms of epistemic states, in which the associated belief set is represented by a formula; for uniformity, we remain with the preceding terminology.

Thus if one revises by $(p \wedge q)$ and then by the negation of some of this information $(\neg p)$, then the other original information (q) is lost. So, in a variant of an example from [DP97], consider where I see a bird in the distance and come to believe that it is red and flies. If on closer examination I see that it is yellow, then according to $(C2')$ I no longer believe that it flies. Hence this is too strong a condition to reasonably adopt, at least for every revision function in all circumstances. Moreover, for approaches based on [Spo88], such as [DP97], it is not at all obvious how such a result can be avoided.

There has also been work on specific revision operators based on the *distance* between models of a knowledge base and a sentence to be incorporated in the knowledge base. This work includes [Bor85, Web86, Dal88, Sat88, Win88, For89]. In these approaches, the models of the new knowledge base are those models of the sentence to be added that are closest (based on "distance" between atomic sentences) to models of the original knowledge base. For example, in [Dal88] the revision operator uses the Hamming distance between interpretations as metric, where the Hamming distance $d(w_1, w_2)$ between interpretations w_1 and w_2 is the number of propositional variables on which the interpretations differ. The distance between an interpretation w and the models of K is given by: $d(Mod(K), w) = \min_{w_i \models K} d(w_i, w)$, where $Mod(K)$ is the set of models of K and $w_i \models K$ indicates that K is true in w_i. A total pre-order on interpretations is given by:

$$w_1 \leq_K w_2 \quad \text{iff} \quad d(Mod(K), w_1) \leq d(Mod(K), w_2).$$

The operator $\dot{+}_D$, defined by $Mod(K \dot{+}_D \alpha) = \min_{\leq_K} Mod(\alpha)$, satisfies the AGM postulates.

[dV93] provides syntactic characterisations of most of the above-cited distance-based approaches. As well, an algorithm is provided for each characterisation. The general strategy is to first convert (a portion of) a knowledge base and formula into disjunctive normal form (DNF). A distance is defined between the clauses in the DNF representations, depending on the approach being considered. Dependencies are propagated among the clauses, generating the set of clauses in the resulting knowledge base. In related work, [EG92] considers the decision problem "*Is p true in $K \dot{+} q$?*" for a wide selection of distance-based operators, and syntactic restrictions on K, q, and p. [LS97] considers how distance-based operators operators can be expressed using circumscription (and vice versa) along with the complexity of the reductions.

A separate direction in belief revision is to assume that revision is not carried out on a belief set per se, but rather on an arbitrary set of formulas. This notion of *base revision* is proposed in [Mak85, FUV83], and fully explored in [Neb92]. The idea is that a knowledge base is represented by a (arbitrary, syntactic) *belief base* that is to be modified, queried, etc. While conceptually simple, revision in these approaches frequently relies on arbitrary syntactic distinctions. With respect to implementations, [Wil95] provides a computational model for belief base revision; other relevant work includes [BDP01] and [Lib99]. These approaches are further discussed and compared with the present approach in Section 6.

Revision and contraction reflect the intuition that an agent receives new information concerning a static world or domain. [KM92] explores the distinct notions of belief *update* and *erasure* in which an agent changes its beliefs in response to changes in its external environment. As well, recently there has been significant interest in belief *merging* or *fusing*, where two or more knowledge sources are combined. Our interests in this paper centre on revision and contraction; as will become apparent, the present approach can be easily extended to represent these other operations.

3 Consistency-Based Belief Change

This section informally introduces our approach to belief change, concentrating on belief revision. As well as describing underlying intuitions and the approach, we also discuss the broader paradigm of consistency-based reasoning.

3.1 A Naïve Approach

The problem we address is the general problem of belief revision:

> Given a general knowledge base and sentence for revision (contraction, etc.), what should the revised (contracted, etc.) knowledge base look like?

A common assumption is that K is to be minimally changed, in order to accommodate α. In our approach, we require that α is true in $K\dot{+}\alpha$, and we subsequently "add" whatever we can from K.

An obvious way to realise such a scheme is to consider an enumeration of sentences of K and, beginning with α, iteratively add each sentence to a candidate revision whenever consistent. Let $\langle\phi_i\rangle_{i\in I}$ be an exhaustive enumeration of the sentences of belief set K, and let α be the sentence for revision. Define:

1. $K_0 = \alpha$.
2. If $K_i \cup \{\phi_i\} \nvdash \bot$
   (a) then $K_{i+1} = Cn(K_i \cup \{\phi_i\})$
   (b) otherwise $K_{i+1} = Cn(K_i)$.

Define $K\dot{+}_I\alpha$ as $\bigcup_{i\in I} K_i$ and $K\dot{+}\alpha$ as $\bigcap_I K\dot{+}_I\alpha$ over all enumerations $\langle\phi_i\rangle_{i\in I}$ of K.

Theorem 3.1 *Let K be a belief set and α a formula such that $K \vdash \neg\alpha$ and $\alpha \nvdash \bot$.*

1. *For every $\beta \in \mathcal{L}$ where $\alpha \nvdash \neg\beta$, there is an enumeration $\langle\phi_i\rangle_{i\in I}$ of K such that $K\dot{+}_I\alpha \vdash \beta$.*
2. *For every enumeration $\langle\phi_i\rangle_{i\in I}$ of K and for every formula β, we have that*

$$K\dot{+}_I\alpha \vdash \beta \quad \textit{or} \quad K\dot{+}_I\alpha \vdash \neg\beta\ .$$

3. $K\dot{+}\alpha = Cn(\alpha)$.

Proof 3.1 (Outline) The proofs are straightforward, and follow those in [AM82] showing similar results for full meet and maxichoice belief change. The key step is to note that since $K \vdash \neg\alpha$ and K is a belief set, we also have $K \vdash \neg\alpha \vee \gamma$. Hence the addition of a sentence $\neg\alpha \vee \gamma$ to a set containing α, in the proposed definition for revision, effectively adds γ. ■

The properties given in Theorem 3.1 are unappealing. Moreover, these difficulties are not easily repaired. For example, in the definition of $K\dot{+}_I\alpha$, if we don't take the deductive closure, via $Cn(\cdot)$, we get the same results. Second, if we just consider enumerations ordered by the logical strength

of formulas, we also get the same results given in Theorem 3.1. Third, if we relax the assumption that K be a belief set, and allow K to be a belief base (i.e. an arbitrary set of formulas), then we essentially obtain the approach to base revision of [FUV83], also explored in [Neb92]. In standard approaches to base revision, among other things, we lose the principle of irrelevance of syntax, given as AGM postulate $(K\dot{+}6)$.

On this last point, [Neb92, p. 58] concludes that abstracting from a syntactic representation of a belief base to a belief set leads nowhere. Nebel goes on to note that several authors (e.g. [Dal88, Win88, KM91]) as a result advocate approaches based on the models characterising a knowledge base and formula. Our approach, introduced informally next, can be seen as a compromise, where a knowledge base and formula can (ultimately) be represented as arbitrary formulas, yet wherein irrelevance of syntax obtains.

3.2 Our Approach

In general, the syntactic form of a sentence doesn't give a clear indication as to which sentences should or should not be retained in a revision. Alternately, one can consider interpretations, and look at the models of K and α. The interesting case occurs when $K \cup \{\alpha\}$ is unsatisfiable because K and α share no models. Intuitively, a model of $K\dot{+}\alpha$ should then contain models of α, but incorporating "parts" of models of K that don't conflict with those of α. That is, we will have

$$Mod(K\dot{+}\alpha) \subseteq Mod(\alpha),$$

and for $m \in Mod(K\dot{+}\alpha)$ we will want to incorporate whatever we can of models of K.

We accomplish this by expressing K and α in different languages, but such that there is an isomorphism between atomic sentences of the languages, and so between the languages themselves. In essence, we replace every occurrence of an atomic sentence p in K by a new atomic sentence p', yielding knowledge base K' and leaving α unchanged. Clearly, under this relabelling, the models of K' and α will be independent, and $K' \cup \{\alpha\}$ will be satisfiable (assuming that each of K, α are satisfiable). We now assert that the languages agree on the truth values of corresponding atoms wherever consistently possible. So, for every atomic sentence p, we assert that $p \equiv p'$ whenever this is consistent with $K' \cup \{\alpha\}$ along with the set of equivalences obtained so far. We obtain a maximal set of such equivalences, call it EQ, such that $K' \cup \{\alpha\} \cup EQ$ is consistent. A model of $K' \cup \{\alpha\} \cup EQ$ then will be a model of α in the original language, wherein the truth values of atomic sentences in K' and α are linked via the set EQ. A candidate "choice" revision of K by α consists of $K' \cup \{\alpha\} \cup EQ$ re-expressed in the original language. General revision corresponds to the intersection of all candidate choice revisions.

To illustrate, consider where

$$K = Cn(\{(p \vee q) \wedge r\}) \qquad \text{and} \qquad \alpha = (\neg p \vee \neg q) \wedge \neg r.$$

Renaming the atoms in K gives $K' = Cn(\{(p' \vee q') \wedge r'\})$. Clearly $K' \cup \{\alpha\}$ is consistent, even though $K \cup \{\alpha\}$ is not. We have that $Cn(K' \cup \{\alpha\} \cup \{p' \equiv p, q' \equiv q\})$ is consistent, but $Cn(K' \cup \{\alpha\} \cup \{p' \equiv p, q' \equiv q, r' \equiv r\})$ is not. Hence we take $EQ = \{p' \equiv p, q' \equiv q\}$. Intersecting $Cn(K' \cup \{\alpha\} \cup EQ)$ with the original language yields $Cn(\{(p \equiv \neg q) \wedge \neg r\})$ as the revised knowledge base.

We can justify this process is as follows: A language has implicit inductive commitments, expressed in the choice of atomic propositions. That is, the atoms are (pragmatically) chosen because they are intended to *mean* something relevant in the domain of discourse. The collection of atomic sentences represents the basic set of meaningful propositions from which further propositions are constructed. In the approach, we essentially employ something resembling a frame assumption, asserting that the truth value of the atomic sentences do not change unless "forced" to change by an incompatibility between K and α. This also means that if we change the representation language, the results of revision may, not unnaturally, change; see [Som94] for a discussion on the sensitivity of revision to the underlying language.

Overall this yields a specific approach to belief revision. The general framework (next section) also allows the expression of contraction and integrity constraints. Further, the general approach also allows the expression of update, erasure, and knowledge base merging operations [DS02]. Significantly, the approach is independent of how the knowledge base and formula for revision are represented. As well, as we show in subsequent sections, the belief change operators have reasonable properties and are well-suited for implementation.

3.3 Consistency-Based Reasoning

The overall approach to belief change described here is founded on the same intuitions as a group of closely-related *consistency-based* reasoning methodologies in Artificial Intelligence. Consistency-based reasoners can be broadly characterised as essentially involving

1. a nonmonotonic minimisation (or maximisation) step that is
2. based on a distinguished set of atoms.

In Theorist [Poo88] for example, one can make predictions of default properties based on selecting from a set of hypotheses, such that the hypotheses selected, together with the background theory and facts, are consistent. Hypotheses are drawn from a designated set of atoms. Similarly, in consistency-based diagnosis [Rei87b], a diagnosis is a conjecture that some minimal set of components are faulty. That a component c_i is faulty, or abnormal, is expressed by a ground formula $Ab(c_i)$, and the assertion that a minimal set of components is faulty is effected by minimizing the set of positive Ab instances. In assumption-based truth maintenance [dK86], explanations are selected from a designated set of atoms.

The emphasis here is slightly different. The maximisation step is applied to *pairs* of corresponding atoms which are asserted to be equivalent. Hence, we do not have a distinguished set of atoms per se to which the maximisation is applied , but rather a designated set of sentences, viz. a set of equivalences between atoms, that is used in the maximization step.

4 Specifying Belief Change Functions

4.1 Formal Foundations

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 initially identified with deductively-closed sets of formulas, or *belief sets*, and are denoted K, K_1, ... 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. Later we relax this restriction and allow knowledge bases to be arbitrary belief bases. 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 $\mathcal{L}_{\mathcal{P}}$ is defined 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 $|C| = 1$.

We next define the notion of an extension for a belief change scenario, called a *belief change extension*. In the definition below, "maximal" is with respect to set containment (rather than set cardinality). The following is our central definition.[2]

Definition 4.1 *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

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

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.*

The sole use of "$\{\bot\}$" in the definition is to take care of the case where $C = \emptyset$. The consistency condition on belief change extensions can be written equivalently as follows:

Alternative Consistency Condition: $K' \cup R \cup EQ \not\vdash \phi$ for every $\phi \in C \cup \{\bot\}$.

We make use of this alternative formulation in the proofs of the theorems.

[2]Our technique of maximizing sets of equivalences of propositional letters bears a superficial resemblance to the use of such equivalences in [LS97] (based in turn on techniques developed in [dK89]). However the approaches are distinct; in particular and in contradistinction to these references, we employ disjoint alphabets for a knowledge base and revising sentence.

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 later use, we define $\overline{EQ}$ as $\{p \equiv p' \mid p \in \mathcal{P}\} \setminus EQ$.

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 and contraction, in Definition 4.2 and 4.3, 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.

The following theorem provides elementary results that will be useful later.

Theorem 4.1 *Let K be a knowledge base and $\alpha \in \mathcal{L}_\mathcal{P}$. Let $EQ, EQ^\star \subseteq \{p \equiv p' \mid p \in \mathcal{P}\}$.*

1. *If EQ determines a consistent belief change extension of (K, R, C), then for $(p \equiv p') \in \overline{EQ}$ there is $\phi \in C \cup \{\bot\}$ such that $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash p \equiv \neg p'$.*
2. *If EQ determines a given consistent belief change extension of $(K, R, \{\alpha\})$, then $K' \cup R \cup \{\neg\alpha\} \cup EQ \vdash p \equiv \neg p'$ for every $(p \equiv p') \in \overline{EQ}$.*
3. *If E_1 and E_2 are two distinct belief change extensions of $(K, R, \{\alpha\})$, then $E_1 \cup E_2 \vdash \bot$.*
4. *If $K \not\vdash \neg\alpha$, then $\{p \equiv p' \mid p \in \mathcal{P}\}$ determines the sole consistent belief change extension of $(K, \{\alpha\}, \emptyset)$.*
5. *If EQ determines a belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$, then EQ determines a belief change extension of $(K, \emptyset, \{\alpha\})$ or of $(K, \emptyset, \{\beta\})$.*
6. *If EQ determines a belief change extension of $(K, \emptyset, \{\alpha\})$, then there is a set of equivalences $EQ^\star$ determining a belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$ such that $EQ \subseteq EQ^\star$.*
7. *EQ determines a belief change extension E_1 of $(K, \{\alpha\}, \emptyset)$ iff EQ determines a belief change extension E_2 of $(K, \emptyset, \{\neg\alpha\})$. Furthermore, $E_1 = Cn(E_2 \cup \{\alpha\})$.*

Parts 1 and 2 of the theorem state that a belief change extension determines the relation between all corresponding pairs of atoms in $\mathcal{P}$ and $\mathcal{P}'$. Part 3 asserts that distinct belief change extensions are mutually inconsistent. The fourth part states that if α is consistent with K then all corresponding atoms in $\mathcal{P}$ and $\mathcal{P}'$ share the same truth value in a (in fact, *the*) resulting belief change extension. The next two parts relate the components of a conjunction comprising C to the individual conjuncts; via Part 7 we get an analogous relation between parts of a disjunction of a formula comprising R. Part 7 of the theorem shows the relation of singleton elements of R and C, along with their respective belief change extensions.

4.2 Revision and Contraction

Definition 4.1 provides a very general framework for specifying belief change. In this subsection we restrict the definition to obtain specific functions for belief revision and contraction. In the definitions below, note that K need not be a belief set, but rather may be any arbitrary set of formulas.

Definition 4.2 (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.*

Observe that for each belief change extension E_i there is some selection function c such that $E_i = K\dot{+}_c\alpha$ and vice versa. A choice revision represents one feasible way in which a knowledge base can be revised to incorporate new information. The intersection of all belief change extensions (comprising skeptical revision) represents a "safe" means of taking all choice revisions into account. One might also take the intersection of *some* set of belief change extensions as the revision of K by α. For example, one may have background information indicating that there is a *preferred* subset of the belief change extensions whose intersection could comprise the revision of K by α. However, we do not address this intermediate notion, analogous to *partial meet* belief change [AGM85].

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.

In detail, for the last example, we wish to determine

$$\{p \wedge q\}\dot{+}(\neg p \vee \neg q)\ . \tag{1}$$

We find determining maximal sets $EQ \subseteq \{p \equiv p', q \equiv q'\}$ such that

$$\{p' \wedge q'\} \cup \{\neg p \vee \neg q\} \cup EQ$$

is consistent. These are: $EQ_1 = \{p \equiv p'\}$ and $EQ_2 = \{q \equiv q'\}$. Accordingly, we obtain

$$\begin{aligned} \{p \wedge q\}\dot{+}(\neg p \vee \neg q) &= \bigcap_{i=1,2} Cn(\{p' \wedge q'\} \cup \{\neg p \vee \neg q\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P} \\ &= Cn(p \equiv \neg q)\,. \end{aligned}$$

In this example there are two choice extensions, $Cn(p \wedge \neg q)$ and $Cn(\neg p \wedge q)$. This raises the question of the usefulness of choice revision compared to general revision. A choice reasoner may be expected to be faster than a full, skeptical, reasoner, since only one extension is generated. However the conclusions obtained from a single extension may be overly strong, since they won't be tempered by those in other extensions. In belief revision this may be less of a problem than, say, in nonmonotonic reasoning: the goal in revision is to determine the true state of the world; if a (choice) revision results in an inaccurate knowledge base, then *this* inaccuracy will presumably be detected and rectified in a later revision. So, over several revisions, choice revision may converge to the true state of the world as quickly as skeptical revision. Hence for a land vehicle exploring a benign environment, choice revision might be an effective part of a control mechanism; for something like flight control, or controlling a nuclear reactor, one would prefer the more conservative skeptical revision.

Contraction is defined similarly to revision.

Definition 4.3 (Contraction) *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 $(K, \emptyset, \{\alpha\})$. Then, we define*

1. $K\dot{-}_c\alpha = E_i$ *as a* choice contraction *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) contraction *of K by α.*

A choice contraction represents a feasible way in which a knowledge base can be contracted to incorporate new information, while the intersection of all choice contractions represents a "safe," skeptical means of taking all choice contractions into account.

Table 2 gives examples of skeptical contraction, using the same format as Table 1. For the

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

Table 2: Skeptical contraction examples.

first example we wish to determine $\{p \wedge q\}\dot{-}q$. To compute the belief change extensions of $(\{p \wedge q\}, \emptyset, \{q\})$ we rename the propositions in $\{p \wedge q\}$ and look for maximal subsets EQ of $\{p \equiv p', q \equiv q'\}$ such that $\{p' \wedge q'\} \cup \{\neg q\} \cup EQ$ is consistent. Thus $EQ = \{p \equiv p'\}$, yielding

$$\begin{aligned} \{p \wedge q\}\dot{-}q &= Cn(\{p' \wedge q'\} \cup \emptyset \cup \{p \equiv p'\}) \cap \mathcal{L}_\mathcal{P} \\ &= Cn(\{p\})\,. \end{aligned}$$

We thus get p, along with all of its logical consequences.

The general approach, with $|C| > 1$, can be immediately employed to express *multiple contraction* [Fuh88], in which contraction is with respect to a set of (not necessarily mutually consistent) sentences. Hence we can use a belief change scenario of the form $(K, \emptyset, \{\alpha, \neg\alpha\})$ to represent a (say) *symmetric contraction* [KM92] of α from K. See Section 4.4 for a related discussion.

4.3 Properties of Revision and Contraction

With respect to the AGM postulates, we obtain the following.[3]

Theorem 4.2 *Let $\dot{+}$ and $\dot{+}_c$ be given as in Definition 4.2.*
Then $\dot{+}$ and $\dot{+}_c$ satisfy the following postulates.[4]

1. $(K\dot{+}1)$ *to* $(K\dot{+}4)$, $(K\dot{+}6)$, $(K\dot{+}7)$
2. $(K\dot{+}5)'$ $K\dot{+}\alpha = K_\perp$ *iff:* $K = K_\perp$ *or* $\vdash \neg\alpha$ *(a weaker version of* $(K\dot{+}5)$*).*
3. $(K\dot{+}6)'$ *If* $\vdash K_1 \equiv K_2$ *and* $\vdash \alpha \equiv \beta$ *then* $K_1\dot{+}\alpha = K_2\dot{+}\beta$ *(a stronger version of* $(K\dot{+}6)$*).*

Hence the basic AGM postulates are (effectively) satisfied, while one of the two supplementary postulates is not. The following is a counterexample to $(K\dot{+}8)$ [KM91, p. 272]:

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

So $(K\dot{+}\alpha)+\beta$ is $(p \wedge \neg q \wedge \neg r \wedge \neg s)$ while $K\dot{+}(\alpha \wedge \beta)$ is $(\neg p \wedge \neg q \wedge r \wedge s) \vee (p \wedge \neg q \wedge \neg r \wedge \neg s)$.

We obtain analogous results for $\dot{-}$ and $\dot{-}_c$ with respect to the AGM contraction postulates:

Theorem 4.3 *Let $\dot{-}$ and $\dot{-}_c$ be given as in Definition 4.3.*
Then, $\dot{-}$ and $\dot{-}_c$ satisfy the following postulates.

1. $(K\dot{-}1)$ *to* $(K\dot{-}3)$, *and* $(K\dot{-}6)$,
2. $(K\dot{-}4)'$ *If* $\vdash K \neq K_\perp$ *and* $\nvdash \phi$ *then* $\phi \notin K\dot{-}\phi$ *(a weaker version of* $(K\dot{-}4)$*).*
3. $(K\dot{-}6)'$ *If* $\vdash K_1 \equiv K_2$ *and* $\vdash \alpha \equiv \beta$ *then* $K_1\dot{-}\alpha = K_2\dot{-}\beta$ *(a stronger version of* $(K\dot{-}6)$*).*

In addition, $\dot{-}$ satisfies the following postulate.

4. $(K\dot{-}7)$ $K\dot{-}\alpha \cap K\dot{-}\beta \subseteq K\dot{-}(\alpha \wedge \beta)$.

[3] If K_1 and K_2 are sets of formulas, we take $\vdash K_1 \equiv K_2$ to mean that $K_1 \vdash \alpha$, $\forall \alpha \in K_2$, and vice versa.

[4] In Definitions 4.1 and 4.2, we have given what seems to us to be the most natural approach to (consistency-based) revision. These definitions yield a slightly weaker version of $(K\dot{+}5)$. To obtain $(K\dot{+}5)$ one can either modify Definition 4.1 so that when B is inconsistent the belief change extension consists of just the closure of R, or, as in [KM91], simply assume that K is consistent.

For $\dot{-}_c$, we have the following results, corresponding to AGM postulates $(K\dot{-}7)$ and $(K\dot{-}8)$.

Theorem 4.4 *For any selection function c, there is selection function c' such that*

1. $K\dot{-}_c(\alpha \wedge \beta) = K\dot{-}_{c'}\alpha$ *or* $K\dot{-}_c(\alpha \wedge \beta) = K\dot{-}_{c'}\beta$
2. *If* $K\dot{-}_c(\alpha \wedge \beta) \nvdash \neg\alpha$ *then* $K\dot{-}_c(\alpha \wedge \beta) = K\dot{-}_{c'}\alpha$.

The controversial recovery postulate $(K\dot{-}5)$ is not satisfied; a counterexample is given by

$$K = Cn(\ p \wedge q\ )\ , \quad \alpha = p \vee q.$$

We obtain $(K\dot{-}\alpha) + \alpha = Cn(p \vee q)$. Hence $p \in K$ but $p \notin (K\dot{-}\alpha) + \alpha$.

We also obtain the following (near) interdefinability results:

Theorem 4.5 (Levi Identity) $K\dot{+}\alpha = (K\dot{-}\neg\alpha) + \alpha$.

Theorem 4.6 (Partial Harper Identity) $K\dot{-}\alpha \subseteq K \cap (K\dot{+}\neg\alpha)$.

The following example shows that equality fails in the Harper Identity: if $K \equiv p \wedge q \wedge r$ and $\alpha \equiv p \wedge q$, then $K\dot{-}\alpha \equiv (p \vee q) \wedge r$ while $K \cap (K\dot{+}\neg\alpha) \equiv (p \equiv \neg q) \wedge r$. Similar results are obtained for choice revision and contraction by appeal to appropriate selection functions.

The operator $\dot{+}$ provides a (near) syntactic counterpart to the minimal-distance-between-models approach of [Sat88]. For two sets S and T, let $S\Delta T$ be the symmetric difference, $(S \cup T) \setminus (S \cap T)$. For formulas α, β, define

$$\Delta^{\min}(\alpha, \beta) \quad = \quad \min_{\subseteq}(\{M_1 \Delta M_2 \mid M_1 \in Mod(\alpha), M_2 \in Mod(\beta)\}),$$

where we identify a model with the set of literals true in the model. Then, we have:

Theorem 4.7 *Let $B = (K, R, \emptyset)$ be a belief change scenario in $\mathcal{L}_\mathcal{P}$ where $K \neq K_\perp$, and let $(EQ_i)_{i \in I}$ be the family of all sets of equivalences, as given in Definition 4.1.*

Then, we have $\{\ \{p \in \mathcal{P} \mid (p \equiv p') \notin EQ_i\}\ \mid i \in I\} \ = \ \Delta^{\min}(K, R)$.

Corollary 4.8 *For any K and α, $K\dot{+}\alpha = K\dot{+}_s\alpha$ where $\dot{+}_s$ is the Satoh revision operator.*

This correspondence provides a semantics for a restriction (viz. skeptical revision) of our general approach. However, we emphasise that the approaches are distinct. First, contraction is expressed here in terms of belief change scenarios, a topic not addressed in Satoh's or other distance-based approaches. Theorem 4.6 shows that contraction can't simply be introduced via the Harper Identity without violating Definition 4.1. As we show in the next section, the implementation of contraction is quite different from that of revision. Lastly, the choice approach, "joint" revision and contraction, and (below) integrity constraints, are not readily expressed in distance-based semantics. [5]

Since we can determine a revision for *every* K and α, the approach clearly supports iterated revision. Indeed, there are nontrivial results concerning iterated revision that hold for the present approach. For example,[6] we have:

[5] An analogy may be drawn to Theorist [Poo88] or the causal calculator [GLMT02]. These approaches begin from independent intuitions, yet are expressible by fragments of default logic or extended logic programs, respectively.

[6] This theorem relies on that fact that we can express knowledge bases and the results of revision as formulas; this is covered in the next section.

Theorem 4.9 *Let $\dot{+}$ be defined as in Definition 4.2. Then, we have*

1. $(\alpha\dot{+}\beta)\dot{+}\alpha = \beta\dot{+}\alpha$.
2. $\beta\dot{+}(\beta\dot{+}\alpha) = \beta\dot{+}\alpha$.
3. $(\alpha\dot{+}\beta)\dot{+}\alpha = \alpha\dot{+}(\beta\dot{+}\alpha)$.

A revision $\alpha\dot{+}\beta$ is often interpreted as comprising that part of β that in some sense is "closest" or "most similar to" the knowledge base given by α. Under this reading, $(\alpha\dot{+}\beta)\dot{+}\alpha$ is the revision of that part of β that is closest to α, by α; Part 1 of the theorem then says that this revision is the same as $\beta\dot{+}\alpha$. In other words, the part of β that plays a role in the revision $\beta\dot{+}\alpha$ is given by $\alpha\dot{+}\beta$. Theorem 4.9.2 has an analogous reading, that the part of α that plays a role in the revision $\beta\dot{+}\alpha$ is exactly given by $\beta\dot{+}\alpha$. Combining Theorem 4.9.1 with the simple result $\beta\dot{+}\alpha = \alpha\dot{+}(\beta\dot{+}\alpha)$ yields Theorem 4.9.3. See [DS02] for a further discussion of iterated revision in this framework.

4.4 Integrity Constraints

Definition 4.1 allows simultaneous revision and contraction by sets of formulas. This in turn 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 [Kow78, SK87], K satisfies IC iff $K \cup \{IC\}$ is satisfiable. In the *entailment-based* approach of [Rei84], K satisfies IC iff $K \vdash IC$. Neither definition is wholly satisfactory; as well, there are others [Rei87a]. [KM91] shows how entailment-based integrity constraints can be maintained across revisions: given an integrity constraint IC (represented as a propositional formula) and revision function $\dot{+}$, a revision function $\dot{+}^{IC}$ which preserves IC is defined by: $K\dot{+}^{IC}\alpha = K\dot{+}(\alpha\wedge IC)$. In our approach, we can define revision taking into account both approaches to integrity constraints.

Corresponding to Definition 4.2 (and ignoring the choice approach) we obtain:

Definition 4.4 *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).*

[SK87] assumes that the set of consistency-based integrity constraints is mutually consistent; in our approach this would correspond to considering belief change scenario

$$(K, \{\alpha\} \cup IC_e, \{\neg \textstyle\bigwedge_{\phi\in IC_c} \phi\}) \quad \text{instead of} \quad (K, \{\alpha\} \cup IC_e, \overline{IC_c}).$$

That is, in our approach, elements of IC_c are individually consistent with respect to a belief change extension. This permits for example $IC_c = \{p, \neg p\}$ to be a nontrivial set of consistency-based integrity constraints (in which the resulting knowledge base remains uncommitted with regards the truth value of p). The next theorem shows that integrity constraints preserve their respective forms of integrity.

Theorem 4.10 *Let $\dot{+}^{(IC_e,IC_c)}$ be defined as in Definition 4.4. Then, we have*

1. $\left(K\dot{+}^{(IC_e,IC_c)}\alpha\right) \vdash IC_e$.
2. *If $K \nvdash \bot$ then:*

   for every $\gamma \in IC_c$: if we have $IC_e \cup \{\alpha\} \nvdash \neg\gamma$ then $\left(K\dot{+}^{(IC_e,IC_c)}\alpha\right) \nvdash \neg\gamma$.

Finally, and in contrast with previous approaches, it is straightforward to add *dynamic* integrity constraints, which express constraints that hold between states of the knowledge base before and after revision. The simplest way of so doing is to add the negation of such constraints to the set C in Definition 4.1. To state that if $a \wedge b$ is true in a knowledge base before revision then c must be true afterwards, we would add $\neg(a' \wedge b' \supset c)$ to C. Note however that the addition of dynamic constraints may lead to an operator that violates some of the properties of $\dot{+}$. For example $Cn(\alpha)\dot{+}\neg\alpha$ with dynamic constraint $\alpha' \supset \alpha$ leads to an inconsistent revision.

5 Implementability Considerations

In this section we address general implementability issues. First we consider the problem of representing the results of revision in a finite, manageable representation. Second, we address limiting the range of EQ. Following this we present a high-level algorithm for implementing the approach; as well we show how the approach can be expressed in Default Logic. Two specific implementations are briefly reviewed, and we finish by giving several complexity results.

5.1 Finite Representations

Definitions 4.1, 4.2, and 4.3 provide a characterisation of revision and contraction, yielding in each case a deductively-closed belief set. Here we consider how the same (with respect to logical equivalence) operators can be defined so that they yield a knowledge base consisting of a (finite) formula. It proves to be the case that, for formulas K and α, we can define choice revision so that the size of $K\dot{+}_c\alpha$ is no greater than the sum of the sizes of K and α for any selection function c.

Informally the procedure is straightforward, although the technical details are less so. A knowledge base K is now represented by a formula. For simplicity we lightly abuse notation in this section, and allow the first argument of a belief change scenario to also be a single formula. Whether a single formula or set of formulas is intended will be clear from the context.

Via Definitions 4.1 and 4.2 we consider maximal sets EQ where $\{K'\}\cup\{\alpha\}\cup EQ$ is consistent. For each such set EQ, we carry out the substitutions:

- for $p \equiv p' \in EQ$, substitute p uniformly for p' in K',
- for $p \equiv p' \notin EQ$, substitute $\neg p$ uniformly for p' in K'.

The result of these substitutions into $K' \wedge \alpha$ is a sentence of size $\leq |K| + |\alpha|$ in language $\mathcal{L}_\mathcal{P}$ and whose deductive closure is equivalent to (some) choice revision. The disjunction of all such sentences (and so considering all possible sets EQ) is equivalent to $K\dot{+}\alpha$.

Observe that any set of equivalences EQ induces a binary partition of its underlying alphabet $\mathcal{P}$, namely $\langle \mathcal{P}_{EQ}, \mathcal{P}_{\overline{EQ}} \rangle$ with $\mathcal{P}_{EQ} = \{p \in \mathcal{P} \mid p \equiv p' \in EQ\}$ and $\mathcal{P}_{\overline{EQ}} = \mathcal{P} \setminus \mathcal{P}_{EQ}$. Given a belief change scenario B along with a set of (determining) equivalences EQ_i (according to Definition 4.1), we define for $\phi \in \mathcal{L}_\mathcal{P}$, that $\lceil \phi \rceil_i$ is the result of replacing in ϕ each proposition $p \in \mathcal{P}_{\overline{EQ_i}}$ by its negation $\neg p$.

Definition 5.1 *Let $B = (K, R, C)$ be a belief change scenario in $\mathcal{L}_\mathcal{P}$ and let $(EQ_i)_{i \in I}$ be the family of all sets of equivalences, as defined in Definition 4.1. Then, we define*

1. $\lceil B \rceil_c$ *as* $\lceil K \rceil_k$ *for some selection function c with $c(I) = k$.*
2. $\lceil B \rceil$ *as* $\bigvee_{i \in I} \lceil K \rceil_i$.

Accordingly, we define

1. $\lceil (K, \{\alpha\}, \emptyset) \rceil_c \wedge \alpha$ as the finite representation of $K\dot{+}_c\alpha$ and
2. $\lceil (K, \{\alpha\}, \emptyset) \rceil \wedge \alpha$ as the finite representation of $K\dot{+}\alpha$.

We have the following result.

Theorem 5.1 *Let $K, \alpha \in \mathcal{L}_\mathcal{P}$. Then, for $(EQ_i)_{i \in I}$ as given in Definition 4.1, we have*

$$K\dot{+}\alpha \quad \equiv \quad \lceil (K, \{\alpha\}, \emptyset) \rceil \wedge \alpha \quad = \quad \bigvee_{i \in I} \lceil K \rceil_i \wedge \alpha.$$

Consider $\{p \wedge q\}\dot{+}(\neg p \vee \neg q)$. So $B = (\{p \wedge q\}, \{\neg p \vee \neg q\}, \emptyset)$. We obtain:

$$\lceil B \rceil \wedge (\neg p \vee \neg q) \qquad = \qquad [(p \wedge \neg q) \vee (\neg p \wedge q)] \wedge (\neg p \vee \neg q) \qquad \equiv \qquad (p \equiv \neg q).$$

Contraction is handled somewhat differently. This is not surprising, given that revision and contraction are not fully interdefinable (Theorem 4.6). In revision we replace each atomic proposition in $\overline{EQ_i}$ by its negation in K. For contraction, we need to substitute into K all possible combinations of truth value assignments for all elements in $\overline{EQ_i}$. As [Lin00] points out, this notion of "forgetting" was first defined by Boole in 1854; it has reappeared in [Web86, LR94, Lin00].

Given a belief change scenario B, a set of equivalences EQ_i (according to Definition 4.1) along with its induced partition $\langle \mathcal{P}_{EQ_i}, \mathcal{P}_{\overline{EQ_i}} \rangle$ of $\mathcal{P}$, we consider the set of functions

$$\Pi_i = \{\pi_{i_k} \mid \pi_{i_k} : \mathcal{P}_{\overline{EQ_i}} \to \{\top, \bot\}\}.$$

For each $\pi_{i_k} \in \Pi_i$ and $\phi \in \mathcal{L}_\mathcal{P}$, we define $\lfloor \phi \rfloor^{i_k}$ as the result of replacing in ϕ each proposition $p \in \mathcal{P}_{\overline{EQ_i}}$ by $\pi_{i_k}(p)$. Note that every set of equivalences EQ_i induces a whole set Π_i of such mappings π_{i_k}, amounting to all possible truth assignments to $\mathcal{P}_{\overline{EQ_i}}$.

Definition 5.2 *Let B and $(EQ_i)_{i\in I}$ be defined as in Definition 4.1. Then, we define*

1. $\lfloor B \rfloor_c$ *as* $\bigvee_{\pi_j \in \Pi_k} \lfloor K \rfloor_k^j$ *for some selection function c with $c(I) = k$.*
2. $\lfloor B \rfloor$ *as* $\bigvee_{i\in I, \pi_j \in \Pi_i} \lfloor K \rfloor_i^j$.

Accordingly, we define

1. $\lfloor (K, \emptyset, \{\alpha\}) \rfloor_c$ as the finite representation of $K \dot{-}_c \alpha$ and
2. $\lfloor (K, \emptyset, \{\alpha\}) \rfloor$ as the finite representation of $K \dot{-} \alpha$.

We have the following result.

Theorem 5.2 *Let $K, \alpha \in \mathcal{L}_\mathcal{P}$. Then, for $(EQ_i)_{i\in I}$ as given in Definition 4.1, we have*

$$K \dot{-} \alpha \quad \equiv \quad \lfloor (K, \emptyset, \{\alpha\}) \rfloor \quad = \quad \bigvee_{i\in I, \pi_j \in \Pi_i} \lfloor K \rfloor_i^j \, .$$

Consider $(p \wedge q) \dot{-} q$. We obtain

$$\lfloor \, (\{p \wedge q\}, \emptyset, \{q\}) \, \rfloor \; = \; (p \wedge \bot) \vee (p \wedge \top) \; \equiv \; p \, .$$

Theorems 5.1 and 5.2 show that revision and contraction can be defined with respect to syntactic objects (viz. a formula for K) yet are essentially independent of syntactic form. That is, whether a knowledge base is represented by a formula, or a set of formulas, if $K_1 \equiv K_2$ and $\alpha_1 \equiv \alpha_2$ then $K_1 \dot{+} \alpha_1 \equiv K_2 \dot{+} \alpha_2$ (and similarly for contraction). Hence in a certain sense the approach combines the advantages of base revision [Neb92] and syntax-independent approaches: knowledge bases and formulas can be represented arbitrarily, yet the results of belief change are independent of syntactic form.

5.2 Limiting the range of EQ

Intuitively, if an atomic sentence appears in a knowledge base K but not in the sentence for revision α, or vice versa, then that atomic sentence plays no part in the revision process. This is indeed the case here. In the following, we show that for computing a belief change extension of belief change scenario $B = (K, R, C)$, we need consider just those atoms common to K and to $R \cup C$.[7]

Let $\mathcal{P}(\phi)$ be the atomic sentences in formula, or set of formulas, ϕ. Recall the notation: for $\alpha \in \mathcal{L}_\mathcal{P}$, the formula α' is obtained by replacing every atomic sentence p in α by p'. This is extended to: for $\mathcal{Q} \subseteq \mathcal{P}$, the formula $\alpha'[\mathcal{Q}]$ is the same as α except that for every $p \in \mathcal{Q}$, where α has p, $\alpha'[\mathcal{Q}]$ has p'. This notation is extended to sets of formulas in the expected fashion. Definition 4.1 is modified to apply to a restricted set of atoms:

[7] In a related but orthogonal vein, [dV93, Par99] split a knowledge base into (effectively) relevant and irrelevant parts. Such techniques could also be used to improve an implementation. We don't pursue the matter here; however see Section 6 for a discussion.

Definition 5.3 *Let $B = (K, R, C)$ be a belief change scenario in $\mathcal{L}_\mathcal{P}$ and let $\mathcal{Q} \subseteq \mathcal{P}$.*
Define $EQ^\mathcal{Q}$ as a maximal set of equivalences $EQ^\mathcal{Q} \subseteq \{p \equiv p' \mid p \in \mathcal{Q}\}$ such that

$$Cn\big(K'[\mathcal{Q}] \cup R \cup EQ^\mathcal{Q}\big) \cap (C \cup \{\bot\}) \;=\; \emptyset.$$

Then

$$Cn\big(K'[\mathcal{Q}] \cup R \cup EQ^\mathcal{Q}\big) \cap \mathcal{L}_\mathcal{P}$$

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

Similarly we define vocabulary-restricted revision:

Definition 5.4 (Vocabulary-Restricted Revision) *Let K be a knowledge base, α a formula, and $\mathcal{Q} \subseteq \mathcal{P}$. Let $(E_i)_{i\in I}$ be the family of all consistent belief change extensions of $(K, \{\alpha\}, \emptyset)$ with respect to $\mathcal{Q}$. Then, we define*

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

Vocabulary-restricted contraction ($\dot{-}_c^\mathcal{Q}$ and $\dot{-}^\mathcal{Q}$) is defined in the obvious analogous fashion.

The next result shows that one obtains the same belief change extensions if the "context" of change is restricted to atoms common to K and $R \cup C$.

Theorem 5.3 *Let $K \subseteq \mathcal{L}_\mathcal{P}$ and $\alpha \in \mathcal{L}_\mathcal{P}$. Let $\mathcal{Q} = \mathcal{P}(K) \cap \mathcal{P}(\alpha)$. Then, we have*

1. $K \dot{+} \alpha \equiv K \dot{+}^\mathcal{Q} \alpha$.
2. $K \dot{-} \alpha \equiv K \dot{-}^\mathcal{Q} \alpha$.

So for belief change, we need consider just the atomic sentences common to K and to α; we can ignore (with regards EQ) other atomic sentences.

We can combine Theorems 5.1 and 5.3 in the obvious fashion to obtain a finite, vocabulary-restricted formulation of revision that is equivalent to the original. We extend our previous notation as follows: Given a belief change scenario B and for $\mathcal{Q} \subseteq \mathcal{P}$, let EQ_i be a set of (determining) equivalences based on $\mathcal{Q}$ (according to Definition 5.3). Define for $\phi \in \mathcal{L}_\mathcal{P}$, that $\lceil \phi \rceil_i^\mathcal{Q}$ is the result of replacing in ϕ each proposition $p \in \mathcal{Q}_{\overline{EQ_i}}$ by its negation $\neg p$.

Definition 5.5 *Let $B = (K, R, \emptyset)$ be a belief change scenario and let $(EQ_i)_{i\in I}$ be the family of all sets of equivalences with respect to $\mathcal{Q} = \mathcal{P}(K) \cap \mathcal{P}(R)$, as given in Definition 5.3. Then, define*

1. $\lceil B \rceil_c^\mathcal{Q}$ *as* $\lceil K \rceil_k^\mathcal{Q}$ *for some selection function c with $c(I) = k$.*
2. $\lceil B \rceil^\mathcal{Q}$ *as* $\bigvee_{i\in I} \lceil K \rceil_i^\mathcal{Q}$.

Accordingly, we define

1. $\lceil (K, \{\alpha\}, \emptyset) \rceil_c^{\mathcal{Q}} \wedge \alpha$ as the finite representation of $K \dot{+}_c^{\mathcal{Q}} \alpha$ and
2. $\lceil (K, \{\alpha\}, \emptyset) \rceil^{\mathcal{Q}} \wedge \alpha$ as the finite representation of $K \dot{+}^{\mathcal{Q}} \alpha$.

We have the following result.

Theorem 5.4 *Let $K, \alpha \in \mathcal{L}_{\mathcal{P}}$, and let $\mathcal{Q} = \mathcal{P}(K) \cap \mathcal{P}(\alpha)$. Then,*

$$K \dot{+} \alpha \quad \equiv \quad \lceil (K, \{\alpha\}, \emptyset) \rceil^{\mathcal{Q}} \wedge \alpha \quad = \quad \bigvee_{i \in I} \lceil K \rceil_i^{\mathcal{Q}} \wedge \alpha$$

for $(EQ_i)_{i \in I}$ as given in Definition 5.3.

Consider an extension to example (1): $\{p \wedge q \wedge r\} \dot{+} ((\neg p \vee \neg q) \wedge s)$. We have $\mathcal{Q} = \{p, q\}$ and

$$\begin{aligned} \lceil B \rceil^{\mathcal{Q}} \wedge ((\neg p \vee \neg q) \wedge s) &= r \wedge [(p \wedge \neg q) \vee (\neg p \wedge q)] \wedge (\neg p \vee \neg q) \wedge s \\ &\equiv (p \equiv \neg q) \wedge r \wedge s. \end{aligned}$$

Notably, in determining the revision, the EQ sets are drawn from $\{p, q\}$ only.

A finite, vocabulary-restricted version of contraction, obtained by combining Theorems 5.2 and 5.3 and equivalent to the original, is similarly obtained. We omit the details.

5.3 Algorithm

The results of the previous subsections lead to an algorithm for computing a belief change extension for an arbitrary belief change scenario B. We have:

Function: *BeliefChange*: Compute a belief change extension for given belief change scenario.
Input: Belief change scenario $B = (K, R, C)$
Output: For input B, a formula equivalent to some belief change extension of B.
Using:
Function $Atoms(S)$ – Returns the set of atoms in the set of formulas S.
Function $Prime(S, A)$ – S is a set of formulas; A is a set of atoms.
Returns S, but where every atom $p \in A$ is replaced by p'.
Function $Replace(S, At_1, At_2)$ – S is a set of formulas; At_1, At_2 are individual atoms.
Returns S with every occurrence of At_1 replaced by At_2.

Function body:
1. if $K \vdash \bot$ or $R \vdash \bot$ then return $\bot$.
2. $In := Out := \emptyset$.
3. $At := Atoms(K) \cap (Atoms(R) \cup Atoms(C))$.
4. $K' := Prime(K, At)$
5. for each $a \in At$ do {
5.1 if { for each $\phi \in C \cup \{\bot\}$

we have $K' \cup R \cup \{p \equiv p' \mid p \in In \cup \{a\}\} \not\vdash \phi$ }
5.2 then $In := In \cup \{a\}$
5.3 else $Out := Out \cup \{a\}$ }
6 for each $p \in In$
6.1 $K' := Replace(K', p', p)$.
7 for each $p \in Out$
7.1 $K' := Replace(K', p', \neg p)$.
8 return $\left(\left(\bigwedge_{\alpha \in K'} \alpha \right) \wedge \left(\bigwedge_{\beta \in R} \beta \right) \right)$.

This algorithm allows to generate a belief change extension in nondeterministic polynomial time. In other words, an extension can be computed by a deterministic polynomial Turing machine which uses the answers given by an NP oracle. The oracle is in charge of performing the consistency and entailment checks at 1 and 5.1, which are computations doable in nondeterministic polynomial time. It is clear from the algorithm that only a polynomial number of calls to the oracle are needed (see also Section 5.6). Note that the selection function is left implicit in Line 5; it is realised by the particular order chosen when treating the atoms in At.

5.4 Belief change scenarios and default logic

As pointed out in Section 3.3, our approach falls within the category of consistency-based reasoning methodologies. As we show now, there is an intimate connection between belief change scenarios and default theories in Default Logic [Rei80].[8] The following theorem makes this precise by showing that there is a 1–1 correspondence between the set of consistent belief change extensions of a belief change scenarios and the extensions of a particular default theory.

Theorem 5.5 *Let $B = (K, R, C)$ be a belief change scenario, where $C = \{\phi_1, \ldots, \phi_n\}$.*
Let $(E_i)_{i \in I}$ be the family of all extensions of default theory

$$\left(\left\{ \frac{: \; p \equiv p', \neg\phi_1, \ldots, \neg\phi_n}{p \equiv p'} \;\middle|\; p \in \mathcal{P} \right\}, K' \cup R \right) .$$

Then $(E_i \cap \mathcal{L}_\mathcal{P})_{i \in I}$ is the family of all belief change extensions of B, and vice versa.

Similar (yet unconstrained) default theories were also used in [BS98] for modelling different forms of paraconsistent reasoning.

5.5 Implementations

There are two prototype implementations available for computing the results of belief change operations. First, belief revision and belief contraction operators have been axiomatised by means of quantified Boolean formulas [DSTW01], in that for both the general approach and for specific operators, a quantified Boolean formula is given such that satisfying truth assignments to the free

[8] This section assumes a basic familiarity with Default Logic.

variables correspond to belief change extensions in the original approach. Thus, in this case the problem of determining the results of a belief change operation is reduced to that of satisfiability. This axiomatisation also allows us to identify strict complexity bounds for the considered reasoning tasks described in the next subsection. The results given in Sections 5.1 and 5.2 are implemented as a special module of the reasoning system QUIP [EETW00], a prototype tool for solving various nonmonotonic reasoning tasks based on reductions to QBFs.

The second implementation, called COBA [DHS02], is implemented in Java. The program was originally implemented as a stand-alone application, after which an applet interface was designed that is suitable for testing any belief revision software. The interface allows the user to enter sentences to the knowledge base or the revision list through a text box; then they can simply click a button to perform the revision. The revised knowledge base appears in a preview window, and can be subsequently saved. In this manner, iterated revision can be easily carried out. Results from the program may be displayed without simplification, with (limited) simplification, or in CNF or DNF. The implementation is intended as a proof-of-concept, and there is room for considerable improvement, to be addressed in later work.

The prototype implementations can be accessed from

`http://www.cs.sfu.ca/~cl/software.htm` .

5.6 Complexity

We consider briefly the complexity of several decision problems in general belief change scenarios, as well as restrictions to revision and contraction. Specifically, we deal 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 some 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 some formula ϕ, decide whether ϕ is contained in all consistent belief change extensions of B.

The above general tasks can also be relativised to analogous tasks for revision (called RDEFEXT, RCHOICE, and RSKEPTICAL respectively) and contraction (CDEFEXT, CCHOICE, and CSKEPTICAL). The following complexity results are obtained in [DSTW01], strengthening and extending those discussed in [DS00]:

Theorem 5.6 *We have the following completeness results:*

1. DEFEXT, RDEFEXT, *and* CDEFEXT *are* NP*-complete.*
2. CHOICE, RCHOICE, *and* CCHOICE *are* Σ_2^p*-complete.*
3. SKEPTICAL, RSKEPTICAL, *and* CSKEPTICAL *are* Π_2^p*-complete.*

Informally, the above complexity bounds are the results of two factors. First, propositional satisfiability is NP-complete. To this end, we have not yet addressed restrictions on the syntactic form of K or α; however see [EG92]. The second results from the determination of the sets $(EQ_i)_{i \in I}$. Of considerable heuristic value in this case is the fact that (via Theorem 5.3) we can restrict these sets to the atoms common to K and α.

Note that our algorithm from Section 5.3 allows for deciding the first group of problems, viz. DEFEXT, RDEFEXT, and CDEFEXT; in addition, it provides us with some belief change extension.

6 Related Work

In Section 2 we reviewed the area of belief revision, concentrating on its theory. Here we continue the discussion by comparing our approach with other specific approaches. Previous work on implementing belief change can be divided into two groups, essentially consisting of implementations of non-base revision and of base revision. The former group typically have good formal properties (for example, conforming to the AGM postulates) but with inefficient implementations, while the latter group may violate some pertinent postulate (often syntax-independence), while being expected to perform reasonably well. We survey this work in some detail since we claim that our approach bridges these categories, in that we have good formal properties (in particular syntax-independence) yet an implementation may be expected to perform reasonably.

Approaches that satisfy the AGM postulates (or, for update, KM postulates) generally implement a distance-based approach. For example, [CW94] implements the PMA approach to update [Win88] in a process that mimics the original definition: for each model of the knowledge base, the closest models of the update formula are determined; the union of all such models is the new knowledge base. The resulting algorithm satisfies the KM update postulates. However, representing a knowledge base by its set of models is not going to be a compact, nor intuitive, way of representing a KB in general. The approach also allows entailment-based integrity constraints.

[dV93] provides a syntactic characterisation and algorithm for most of the distance-based approaches to revision and update. The formula to be incorporated is assumed to be in DNF; as well the algorithms rely on a "relevant" portion of the knowledge base (see below) being in DNF. Hence these algorithms may require an exponential time step, and exponential space, that our's do not. Revision or update by formula α is restricted to a "relevant" portion of the knowledge base; this consists of those clauses in the knowledge base sharing atoms with α, call them ψ_0, along with, recursively at Step $i+1$, those clauses sharing atoms with clauses in ψ_i.[9] This is distinct from our approach, where EQ sets are drawn just from those atoms common to the knowledge base and formula for revision. Entailment-based integrity constraints are handled in the following manner: First the revision without integrity constraints is computed. If the integrity constraints are true in the result, the process halts. Otherwise the revision is recomputed with the original formula conjoined with those integrity constraints that didn't follow after the original revision. This process is repeated until all integrity constraints are entailed.

[9] [Par99] does something similar in *splitting* the language of a theory. To incorporate del Val's or Parikh's heuristic in our algorithm of Section 5.3: Prior to Line 4, K would be split into relevant and irrelevant parts, K_{rel} and K_{irr}; K_{rel} would be primed and assigned to K' in Line 4; and in Line 8, K_{irr} would be returned as an additional conjunct.

[Lib99] presents a framework in which revision, update, and merging of knowledge bases may be jointly expressed; contraction and erasure are not considered. (As Section 4.3 shows, one can't just use the Harper Identity to obtain these latter operations.) The operators are expressed in terms of a distance-based semantics, in which the AGM (or KM) postulates are claimed to hold. Update corresponds to Forbus' approach [For89] while revision appears to correspond to Dalal's approach [Dal88].[10] As with [CW94], the output of the system is a set of models.

In the above-cited works, the requirement that the knowledge base be in DNF (or represented by its models) will be impractical for many applications or for large knowledge bases. Often, one would expect a knowledge base to consist of a large number of relatively small-sized assertions, and so be relatively close to conjunctive normal form.

For belief base revision, the earliest work appears to be [FUV83], where a revision consists of the formula for revision together with (the disjunction of) all maximal subsets of the knowledge base that are consistent with the formula for revision; no model theoretic analysis is given.

With respect to implementations, [Wil95] provides a computational model for belief base revision based on *partial entrenchment rankings*. The dynamic behaviour of the system is described by a procedure of *adjustment*. Adjusting a sentence down in the ranking reflects a generalised notion of contraction; adjusting upwards reflects a notion of increased acceptance. The adjustment of one sentence may result in the adjustment of other sentences. The result is an intuitively-appealing model for revising and contracting a finite base of beliefs although, as with other such approaches, there is a syntactic sensitivity to how a ranking is expressed. For example the two rankings

$$\mathbf{B_1}(\phi \wedge \psi) = 6,\ \mathbf{B_1}(\phi \vee \psi) = 8 \qquad \text{and} \qquad \mathbf{B_2}(\phi) = \mathbf{B_2}(\psi) = 6,\ \mathbf{B_2}(\phi \vee \psi) = 8$$

are equivalent, yet a contraction of ϕ in $\mathbf{B_1}$ results in a contraction of ψ (since the formula $\phi \wedge \psi$ is adjusted downwards), while in the second case it does not.

[BDP01] gives a framework in which belief change and fusion are expressed in the context of possibility theory. The authors consider change both with respect to possibilistic belief sets and to possibilistic belief bases. While complexity results and algorithms are not given, the syntactic framework appears suitable for the realisation of a variety of belief change operators.

7 Conclusion

We have presented a general consistency-based framework for belief change, having the same flavour as the consistency-based approaches to diagnosis or default reasoning. The approach centres on the notion of a belief change scenario, consisting of a triple of sets of formulas, $B = (K, R, C)$. Informally, K is a knowledge base that is to be modified so that the formulas in R are contained in (or implied by) the result, and the formulas in C are not. We focus initially on approaches to belief revision, where $|R| = 1$ and $C = \emptyset$, and to belief contraction, in which $R = \emptyset$ and $|C| = 1$. To determine a revision $K\dot{+}\alpha$, the knowledge base K and sentence α are expressed in separate languages. Given this, we syntactically force truth assignments to the atoms in the languages of K and α to coincide insofar as consistently possible. Lastly, we express the resultant

[10] It is suggested that Dalal revision is captured in [Lib99], but not in [LS00].

knowledge base in the original language. There may be more than one way in which this process may be carried out. This gives rise to two notions of revision: a choice notion, in which one such "extension" is used for the revised state, and the intersection of all such extensions.

The approach is amenable for implementation: belief change can be expressed in terms of a finite knowledge base; and the scope of a change operation can be restricted to those propositions common to the knowledge base and sentence. Other considerations, such as splitting the language of the knowledge base, are easily incorporated. We give an algorithm for computing a belief change extension, and show how the approach may be realised in Default Logic. There are two prototype implementations, one using quantified Boolean formulas, and the other providing a Java applet.

A primary contribution of the approach is that we combine theoretical and practical aspects in a single system. Our revision and contraction operators have good formal properties, in particular satisfying the majority of the AGM postulates. Notably, the result of a belief change is independent of the syntactic form of the knowledge base and formula for change. As well, the approach is amenable to implementation. For choice revision, the size of the revised knowledge base is bounded by the sum of the size of the knowledge base and formula for revision. In general revision, the size of a resulting knowledge base depends further on the number of (choice) extensions. This contrasts with previous implementations of non-base approaches, which may require exponential space in a DNF representation or in listing a set of models. Unlike previous approaches, we also consider contraction (along with arbitrary combinations of revision and contraction). Notably, given our assumptions, contraction is not interdefinable with revision, and its implementation must be handled differently from that of revision.

The approach allows for a simple, uniform treatment of integrity constraints, including consistency-based and entailment-based static constraints, as well as dynamic constraints. The approach trivially supports iterated revision, since belief change extensions are defined over all triples of formulas. Although we do not do so here (but see [DS02]), it is straightforward to apply the approach to other belief operations such as update, erasure, and merging.

A Proofs

A.1 Proofs of Section 4

Proof 4.1

1. Let $EQ \subseteq \{p \equiv p' \mid p \in \mathcal{P}\}$ be a set of equivalences determining some consistent belief change extension of (K, R, C).

   Assume that $\overline{EQ} \neq \emptyset$, and let $p \equiv p' \in \overline{EQ}$. By the maximality of EQ, we have that $K' \cup R \cup \{\neg\phi\} \cup EQ \cup \{p \equiv p'\} \vdash \bot$ for some $\phi \in C \cup \{\bot\}$. That is, $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash \neg(p \equiv p')$ or equivalently $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash (p \equiv \neg p')$.

2. This is an obvious consequence of the previous part in which $|C| = 1$.

3. This is an immediate consequence of the previous part: since $E_1 \neq E_2$ we get $EQ_1 \neq EQ_2$, from which the result follows.

4. Any model of $K \cup \{\alpha\}$ over $\mathcal{L}_\mathcal{P}$ can be extended to a model of $K \cup \{\alpha\} \cup EQ$ over $\mathcal{L}_{\mathcal{P}\cup\mathcal{P}'}$, where $EQ = \{p \equiv p' \mid p \in \mathcal{P}\}$. Further, a model of $K \cup \{\alpha\} \cup EQ$ over $\mathcal{L}_{\mathcal{P}\cup\mathcal{P}'}$ is a model of $K' \cup \{\alpha\} \cup EQ$. Since we are given that $K \cup \{\alpha\}$ has a model, and since EQ is the maximum set of equivalences, it is, trivially, the only maximal set of equivalences.

5. Let EQ be a maximal set of equivalences determining a belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$. By definition, $K' \cup EQ \not\vdash \alpha \wedge \beta$. Thus $K' \cup EQ \not\vdash \alpha$ or $K' \cup EQ \not\vdash \beta$. Further

$$K' \cup EQ \cup \{e\} \vdash \alpha \wedge \beta \quad \text{for any} \quad e \in \overline{EQ}. \tag{2}$$

   If $K' \cup EQ \not\vdash \alpha$ then EQ is a maximal (from (2)) set of equivalences determining a belief change extension of $(K, \emptyset, \alpha)$.

   Alternately, $K' \cup EQ \not\vdash \beta$ and an analogous result holds for a belief change extension of $(K, \emptyset, \beta)$.

6. We are given that $K' \cup EQ \not\vdash \alpha$; hence $K' \cup EQ \not\vdash \alpha \wedge \beta$. Clearly EQ can be extended to a maximal set of equivalences $EQ^\star \supseteq EQ$ such that $K' \cup EQ^\star \not\vdash \alpha \wedge \beta$, and either $EQ^\star = \{p \equiv p' \mid p \in \mathcal{P}\}$ or $K' \cup EQ^\star \cup \{e\} \vdash \alpha \wedge \beta$ for every $e \in \overline{EQ^\star}$. In either case, $EQ^\star$ determines a belief change extension of $(K, \emptyset, \alpha \wedge \beta)$.

7. **if part:** Let E_2 be a belief change extension of $(K, \emptyset, \{\neg\alpha\})$ given by $Cn(K' \cup EQ) \cap \mathcal{L}_\mathcal{P}$ where $K' \cup EQ \not\vdash \neg\alpha$ and so $K' \cup EQ \cup \{\alpha\} \not\vdash \bot$.

   Thus $Cn(K' \cup EQ \cup \{\alpha\}) \cap \mathcal{L}_\mathcal{P}$ satisfies the definition of a belief change extension of $(K, \{\alpha\}, \emptyset)$. As well,

$$\begin{aligned} E_1 &= Cn(K' \cup EQ \cup \{\alpha\}) \cap \mathcal{L}_\mathcal{P} \\ &= Cn((Cn(K' \cup EQ) \cap \mathcal{L}_\mathcal{P}) \cup \{\alpha\}) \\ &= Cn(E_2 \cup \{\alpha\}) \,. \end{aligned}$$

   only-if part: Let $E_1 = Cn(K' \cup EQ \cup \{\alpha\}) \cap \mathcal{L}_\mathcal{P}$ be a belief change extension of $(K, \{\alpha\}, \emptyset)$. Thus $K' \cup \{\alpha\} \cup EQ \not\vdash \bot$.

   Hence by Definition 4.1, $E_2 = Cn(K' \cup EQ) \cap \mathcal{L}_\mathcal{P}$ is a belief change extension of belief change scenario $(K, \emptyset, \{\neg\alpha\})$.

   By the same argument as in the if part, we get that $E_1 = Cn(E_2 \cup \{\alpha\})$.

■

Proof 4.2

We just give proofs for $\dot{+}$; those for $\dot{+}_c$ follow as corollaries.

$(K\dot{+}1)$, $(K\dot{+}2)$, and $(K\dot{+}6)'$ are obvious.

For $(K\dot{+}3)$, if $K \vdash \neg\alpha$ then $K + \alpha = \mathcal{L}_\mathcal{P}$ and so $K\dot{+}\alpha \subseteq K + \alpha$.

So assume that $K \not\vdash \neg\alpha$. By Theorem 4.1.4 there is a single consistent belief change extension in which $EQ = \{p \equiv p' \mid p \in \mathcal{P}\}$. It follows that $Cn(K' \cup \{\alpha\} \cup EQ) \cap \mathcal{L}_\mathcal{P} = Cn(K \cup \{\alpha\})$:

$\subseteq$**:** We obtain that $Cn(K' \cup EQ) \cap \mathcal{L}_\mathcal{P} \subseteq Cn(K)$ by virtue of the fact that any model of $K' \cup EQ$ is a model of K; the result then follows immediately.

$\supseteq$**:** We need to show that if, for every $\phi \in \mathcal{L}_\mathcal{P}$, $K \cup \{\alpha\} \vdash \phi$ then $K' \cup \{\alpha\} \cup EQ \vdash \phi$. This is the same as, for every $\phi \in \mathcal{L}_\mathcal{P}$, if $K' \cup \{\alpha\} \cup EQ \nvdash \phi$ then $K \cup \{\alpha\} \nvdash \phi$, or:

If $K' \cup \{\alpha\} \cup EQ \cup \{\neg\phi\} \nvdash \bot$ then $K \cup \{\alpha\} \cup \{\neg\phi\} \nvdash \bot$.

But clearly any model of $K' \cup \{\alpha\} \cup EQ \cup \{\neg\phi\}$ is also a model of $K \cup \{\alpha\} \cup \{\neg\phi\}$, from which our result follows.

Hence $K\dot{+}\alpha = Cn(K \cup \{\alpha\}) = K + \alpha$. This also establishes $(K\dot{+}4)$.

For $(K\dot{+}5)'$, if $K = K_\bot$ or $\vdash \neg\alpha$ then $K\dot{+}\alpha = K_\bot$. Otherwise, $K \neq K_\bot$ and $\nvdash \neg\alpha$, and so $K\dot{+}\alpha \neq K_\bot$ by Definition 4.1.

For $(K\dot{+}7)$, the postulate is trivially satisfied if $(K\dot{+}\alpha) + \beta \vdash \bot$. Consequently assume that $(K\dot{+}\alpha) + \beta \nvdash \bot$.

We must show that $K\dot{+}(\alpha \wedge \beta) \subseteq (K\dot{+}\alpha) + \beta$, or, expanding via Definition 4.2,

$$\left(\bigcap_{i\in I} Cn(K' \cup \{\alpha \wedge \beta\} \cup EQ_i)\right) \cap \mathcal{L}_\mathcal{P} \subseteq Cn\left(\left(\left(\bigcap_{i\in I} Cn(K' \cup \{\alpha\} \cup EQ_i)\right) \cap \mathcal{L}_\mathcal{P}\right) \cup \{\beta\}\right).$$

Assume that

$$\bigcap_{i\in I} Cn(K' \cup \{\alpha \wedge \beta\} \cup EQ_i) \vdash \phi \quad \text{where} \quad \phi \in \mathcal{L}_\mathcal{P}$$

To conclude we need to show that

$$\left(\left(\bigcap_{i\in I} Cn(K' \cup \{\alpha\} \cup EQ_i)\right) \cap \mathcal{L}_\mathcal{P}\right) \cup \{\beta\} \vdash \phi$$

or that

$$(Cn(K' \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P}) \cup \{\beta\} \vdash \phi \tag{3}$$

for every belief change extension of $(K, \{\alpha\}, \emptyset)$.

We make use of the following lemma.

Lemma A.1 *If EQ determines a belief change extension of $(K, \{\alpha \vee \beta\}, \emptyset)$, then EQ determines a belief change extension of $(K, \{\alpha\}, \emptyset)$ or of $(K, \{\beta\}, \emptyset)$.*

Proof A.1 Immediate from Theorem 4.1.7 and Theorem 4.1.5 ∎

Let EQ be a set of equivalences determining some belief change extension of $(K, \{\alpha\}, \emptyset)$ or $(K, \{(\alpha \wedge \beta) \vee (\alpha \wedge \neg\beta)\}, \emptyset)$.

From Lemma A.1 we get that EQ determines some belief change extension of $(K, \{\alpha \wedge \beta\}, \emptyset)$ or $(K, \{\alpha \wedge \neg\beta\}, \emptyset)$.

In the former case we have by assumption that $Cn(K' \cup \{\alpha \wedge \beta\} \cup EQ) \vdash \phi$ and so $Cn(K' \cup \{\alpha\} \cup EQ) \cup \{\beta\} \vdash \phi$ as required.

If this case does not hold, then $Cn(K' \cup \{\alpha\} \cup EQ) \vdash \neg\beta$ and so $Cn(K' \cup \{\alpha\} \cup EQ) \cup \{\beta\} \vdash \bot$, thus trivially $Cn(K' \cup \{\alpha\} \cup EQ) \cup \{\beta\} \vdash \phi$.

■

Proof 4.3

We just give proofs for $\dot{-}$; those for $\dot{-}_c$ follow as corollaries, except as noted in Theorem 4.4.

$(K\dot{-}1)$ and $(K\dot{-}6)'$ are obvious.

For $(K\dot{-}2)$ we need to show that if $\phi \in K\dot{-}\alpha$ then $\phi \in K$. As noted in the proof of $(K\dot{+}3)$, this amounts to showing that if $K \cup \{\neg\phi\} \not\vdash \bot$ then $(K\dot{-}\alpha) \cup \{\neg\phi\} \not\vdash \bot$, or: if $K \cup \{\neg\phi\} \not\vdash \bot$ then $\bigcap_{i\in I}(K' \cup EQ_i) \cup \{\neg\phi\} \not\vdash \bot$.

So let M be a model of $K \cup \{\neg\phi\}$ over the language $\mathcal{L}_\mathcal{P}$. We construct a model M' of $\bigcap_{i\in I}(K' \cup EQ_i) \cup \{\neg\phi\}$ over $\mathcal{L}_{\mathcal{P}\cup\mathcal{P}'}$ by: M' assigns $true$ to $p' \in \mathcal{P}'$ iff M assigns $true$ to $p \in \mathcal{P}$. Obviously then M' is a model of $K' \cup EQ \cup \{\neg\phi\}$ for $EQ = \{p \equiv p' \mid p \in \mathcal{P}\}$, and so M' is a model of $K' \cup EQ_i \cup \{\neg\phi\}$ for every $EQ_i \subseteq EQ$, from which our result follows.

For $(K\dot{-}3)$, if $\alpha \notin K$ then $K \cup \{\neg\alpha\} \not\vdash \bot$; hence $K' \cup \{\neg\alpha\} \cup \{p \equiv p' \mid p \in \mathcal{P}\} \not\vdash \bot$; hence K is the sole consistent belief change extension of $(K, \emptyset, \{\alpha\})$; hence $K\dot{-}\alpha = K$.

For $(K\dot{-}4)'$, assume $K \neq K_\bot$ and $\not\vdash \alpha$. For belief change scenario $(K, \emptyset, \{\alpha\})$ we have $K' \cup \{\neg\alpha\} \not\vdash \bot$; hence there is a maximal set of equivalences EQ (Definition 4.1) such that $K' \cup \{\neg\alpha\} \cup EQ \not\vdash \bot$. Hence $K' \cup EQ \not\vdash \alpha$ and so $K\dot{-}\alpha \not\vdash \alpha$.

For $(K\dot{-}7)$, let:

- $EQ^\alpha_1, \ldots, EQ^\alpha_n$ determine the belief change extensions of $(K, \emptyset, \{\alpha\})$
- $EQ^\beta_1, \ldots, EQ^\beta_m$ determine the belief change extensions of $(K, \emptyset, \{\beta\})$

1. For each $EQ \in \{EQ^\alpha_1, \ldots, EQ^\alpha_n, EQ^\beta_1, \ldots, EQ^\beta_m\}$ there exists $EQ^{\alpha\beta} \supseteq EQ$ that determines a belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$ (Theorem 4.1.6).

2. Also for every $EQ^{\alpha\beta}$ that determines a belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$, we have that $EQ^{\alpha\beta}$ determines a belief change extension of $(K, \emptyset, \{\alpha\})$ or of $(K, \emptyset, \{\beta\})$ (Theorem 4.1.5).

Assume that $K\dot{-}\alpha \vdash \phi$ and $K\dot{-}\beta \vdash \phi$. Hence for every EQ, as given in 1., $K' \cup EQ \vdash \phi$. As well, there is $EQ^{\alpha\beta} \supseteq EQ$ (as specified in 1.) that determines a belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$; and from monotonicity we also have $K' \cup EQ^{\alpha\beta} \vdash \phi$. From 2. we get that every belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$ has a corresponding belief change extension of $(K, \emptyset, \{\alpha\})$ or of $(K, \emptyset, \{\beta\})$. It follows that for every belief change extension of $(K, \emptyset, \{\alpha \wedge \beta\})$ determined by $EQ^\star$ we have $K' \cup EQ^\star \vdash \phi$. Hence $K\dot{-}\alpha \cap K\dot{-}\beta \subseteq K\dot{-}(\alpha \wedge \beta)$. ■

Proof 4.4

1. This is a corollary of Theorem 4.1.5.

2. Assume that $K\dot{-}_c(\alpha \wedge \beta) \not\vdash \neg\alpha$. Thus for some set $EQ^{\alpha\beta}$ determining $(K, \emptyset, \{\alpha \wedge \beta\})$ we have $K' \cup EQ^{\alpha\beta} \not\vdash \alpha$ and so $K' \cup EQ^{\alpha\beta} \cup \{\neg\alpha\} \not\vdash \bot$.

Further if $\overline{EQ^{\alpha\beta}} \neq \emptyset$ then $K' \cup EQ^{\alpha\beta} \cup \{e\} \vdash \alpha \wedge \beta$ for any $e \in \overline{EQ^{\alpha\beta}}$; hence $K' \cup EQ^{\alpha\beta} \cup \{e\} \vdash \alpha$.

So $EQ^{\alpha} = EQ^{\alpha\beta}$ is a maximal set of equivalences determining a belief change extension of $(K, \emptyset, \{\alpha\})$. Hence there is a selection function c' (that chooses EQ^{α}) such that $K\dot{-}_c(\alpha \wedge \beta) = K\dot{-}_{c'}\alpha$.

■

Proof 4.5

We have that E_i is belief change extension of $(K, \emptyset, \{\neg\alpha\})$ iff $Cn(E_i \cup \{\alpha\})$ is a belief change extension of $(K, \{\alpha\}, \emptyset)$ (Theorem 4.1.7).

Let $(E_i)_{i\in I}$ be the family of all consistent belief change extensions of $(K, \{\alpha\}, \emptyset)$. Then

$$\begin{aligned} K\dot{+}\alpha &= \bigcap_{i\in I} Cn(E_i) \quad = \quad \bigcap_{i\in I} Cn(E_i \cup \{\alpha\}) \\ &= \bigcap_{i\in I} Cn(Cn(E_i) \cup \{\alpha\}) \quad = \quad Cn\left(\bigcap_{i\in I} Cn(E_i) \cup \{\alpha\}\right) \\ &= Cn\big((K\dot{-}\alpha) \cup \{\alpha\}\big) \\ &= (K\dot{-}\alpha) + \alpha. \end{aligned}$$

■

Proof 4.6

We need to show the two parts:

1. $K\dot{-}\alpha \subseteq K$.

   This is just ($K\dot{-}2$).

2. $K\dot{-}\alpha \subseteq K\dot{+}\neg\alpha$.

   From Theorem 4.1.7 we get that there is a 1-1 correspondence between every belief change extension E_1 of $(K, \emptyset, \{\alpha\})$ and E_2 of $(K, \{\neg\alpha\}, \emptyset)$, where $E_2 = Cn(E_1 \cup \{\alpha\})$, and so $E_1 \subseteq E_2$.

   Hence, if $(E_{1,i})_{i\in I}$ is the family of all consistent belief change extensions of $(K, \emptyset, \{\alpha\})$ and $(E_{2,i})_{i\in I}$ is the family of all consistent belief change extensions of $(K, \{\alpha\}, \emptyset)$, then

$$K\dot{-}\alpha \;=\; \bigcap_{i\in I} E_{1,i} \quad \subseteq \quad \bigcap_{i\in I} E_{2,i} \;=\; K\dot{+}\neg\alpha.$$

■

Proof 4.7

$\supseteq$: Let $\{p_1, \ldots, p_n\} \in \Delta^{\min}(K, R)$.

So there are models M_1 of K and M_2 of R such that $M_1 \Delta M_2 = \{p_1, \ldots, p_n\}$.

Thus we have:

$$\begin{array}{lll} p \in \{p_1, \ldots, p_n\} & \text{iff:} & M_1 \vdash p \text{ iff } M_2 \nvdash p \\ & \text{iff:} & M_1 \vdash p \text{ iff } M_2 \vdash \neg p. \end{array}$$

Thus for K' over $\mathcal{L}_{\mathcal{P}'}$ there is a model M_1' isomorphic to M_1 such that

$$p \in \{p_1, \ldots, p_n\} \quad \text{iff:} \quad M_1' \vdash p' \text{ iff } M_2 \vdash \neg p.$$

Let M'' be the composition of M_1' and M_2 over language $\mathcal{L}_{\mathcal{P} \cup \mathcal{P}'}$. We obtain:

$$\begin{array}{lll} p \in \{p_1, \ldots, p_n\} & \text{iff:} & M'' \vdash p' \text{ iff } M'' \vdash \neg p \\ & \text{iff:} & M'' \vdash p' \equiv \neg p. \end{array}$$

Hence

$$\begin{array}{lll} p \in \mathcal{P} \setminus \{p_1, \ldots, p_n\} & \text{iff:} & M'' \nvdash p' \equiv \neg p. \\ & \text{iff:} & M'' \vdash p' \equiv p. \end{array}$$

Thus $EQ = \{p \equiv p' \mid p \in \mathcal{P} \setminus \{p_1, \ldots, p_n\}\}$ is a set of equivalences such that $K' \cup R \cup EQ \nvdash \bot$. As well, since $\{p_1, \ldots, p_n\} \in \Delta^{\min}(K, R)$, we get for any $p \notin \{p_1, \ldots, p_n\}$, that $K' \cup R \cup EQ \cup \{p \equiv p'\} \vdash \bot$. Hence EQ is a maximal set of such equivalences, and so determines some consistent belief change extension of B.

$\subseteq$: Let $B = (K, R, \emptyset)$ be a belief change scenario in $\mathcal{L}_{\mathcal{P}}$ where $K \neq K_\bot$ and $R \nvdash \bot$, and let EQ be a set of equivalences as given in Definition 4.1.

Then there is an assignment of truth values to atoms in $\mathcal{P} \cup \mathcal{P}'$ where $K' \cup R \cup EQ \nvdash \bot$.

For any model M'' of $K' \cup R \cup EQ$ we have by definition:

$$p \equiv p' \in EQ \quad \text{iff} \quad M'' \vdash p \equiv p'.$$

For model M'' as above, we define models M_1 and M_2 over $\mathcal{L}_{\mathcal{P}}$ by:

$$M_1 \vdash p \quad \text{iff} \quad M'' \vdash p' \qquad \text{and:} \qquad M_2 \vdash p \quad \text{iff} \quad M'' \vdash p$$

Then:

1. M_1 is a model of K (since M'' is a model of K').
2. M_2 is a model of R (since M'' is a model of R).

3. $M_1 \Delta M_2 \in \Delta^{\min}(K, R)$. (This follows from the maximality of EQ.)

From this it follows that $\{p \in \mathcal{P} \mid (p \equiv p') \notin EQ\} \in \Delta^{\min}(K, R)$.

Since EQ is an arbitrary set of equivalences determining a belief change extension of B, our result follows.

■

Proof 4.9

Notation: In Section 4.1, for $\alpha \in \mathcal{L}_{\mathcal{P}}$, we defined α' as being the same as α but with all atoms replaced by primed counterparts. Here (only) we extend the definition to $\alpha \in \mathcal{L}_{\mathcal{P} \cup \mathcal{P}'}$ in the natural fashion: For $\alpha \in \mathcal{L}_{\mathcal{P} \cup \mathcal{P}'}$, α' is the result of replacing in α each proposition $p \in \mathcal{P}$ by the the corresponding proposition $p' \in \mathcal{P}'$, and replacing each proposition $p' \in \mathcal{P}'$ by the the corresponding proposition $p \in \mathcal{P}$. Hence $\alpha = (\alpha')'$ and for a set of equivalences EQ, we have $EQ = EQ'$.

We assume a finite language for expressing a belief change scenario and we rely on the fact that a belief set in such a case can be finitely represented (see Section 5).

We begin with the following lemma

Lemma A.2 *1. EQ determines a belief change extension of $\alpha \dot{+} \beta$ iff EQ determines a belief change extension of $\beta \dot{+} \alpha$.*

2. EQ determines a belief change extension of $(\alpha \dot{+} \beta) \dot{+} \alpha$ iff EQ determines a belief change extension of $\beta \dot{+} \alpha$.

Proof A.2

1. This follows immediately from Definition 4.1.
2. Let EQ be a maximal set of equivalences determining a belief change extension of $(\alpha \dot{+} \beta) \dot{+} \alpha$.
   Then $(\alpha \dot{+} \beta)' \cup \{\alpha\} \cup EQ \nvdash \bot$.
   So for $(EQ_j)_{j \in J}$ determining the belief change extensions of $\alpha \dot{+} \beta$ we have:
   $\bigcap_{j \in J} (Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j) \cap \mathcal{L}_{\mathcal{P}})' \cup \{\alpha\} \cup EQ \nvdash \bot$.
   $\bigcap_{j \in J} (Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j) \cap \mathcal{L}_{\mathcal{P}})' \cup EQ \nvdash \bot$ (since $\alpha = (\alpha')'$).
   $\bigcap_{j \in J} (Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j \cup EQ) \cap \mathcal{L}_{\mathcal{P}})' \nvdash \bot$ (since $EQ = EQ'$).
   For specific EQ_j we have that $EQ_j \cup EQ \vdash \bot$ iff $EQ_j \neq EQ$. Consequently the above simplifies to:
   $((\{\alpha'\} \cup \{\beta\} \cup EQ) \cap \mathcal{L}_{\mathcal{P}})' \nvdash \bot$.
   Thus $(\{\beta'\} \cup \{\alpha\} \cup EQ) \cap \mathcal{L}_{\mathcal{P}} \nvdash \bot$ from which it follows that EQ determines an extension of $\beta \dot{+} \alpha$.
   Since each step in the preceding can be replaced by an "iff" the result follows.

■

1. Let $(EQ_i)_{i\in I}$ be the family of all sets of equivalences determining extensions of $(\alpha\dot{+}\beta, \{\alpha\}, \emptyset)$ and let $(EQ_j)_{j\in J}$ be the family of all sets of equivalences determining extensions of $(\{\alpha\}, \{\beta\}, \emptyset)$. Then:

$$\begin{aligned}
(\alpha\dot{+}\beta)\dot{+}\alpha &= \bigcap_{i\in I} Cn\big(\{\alpha\dot{+}\beta\}' \cup \{\alpha\} \cup EQ_i\big) \cap \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{i\in I} Cn\left(\left(\bigcap_{j\in J} Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j) \cap \mathcal{L}_{\mathcal{P}}\right)' \cup \{\alpha\} \cup EQ_i\right) \cap \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{i\in I}\bigcap_{j\in J} Cn\big((Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j) \cap \mathcal{L}_{\mathcal{P}})' \cup \{\alpha\} \cup EQ_i\big) \cap \mathcal{L}_{\mathcal{P}} \quad (4)
\end{aligned}$$

(4) is of the form $\bigcap_{i\in I}\bigcap_{j\in J} \Psi_{i,j} \cap \mathcal{L}_{\mathcal{P}}$. From Lemma A.2.2, it follows, for specific i and j appearing in the intersections in (4), that $\Psi_{i,j} = \mathcal{L}_{\mathcal{P}\cup\mathcal{P}'}$ if $EQ_i \neq EQ_j$. Moreover from Lemma A.2.2, it follows that for every distinct EQ_i (as indexed by the first intersection in (4)) there is a EQ_j (indexed in the second intersection in (4)) such that $EQ_i = EQ_j$. Conequently we can simplify (4):

$$\begin{aligned}
(4) &= \bigcap_{j\in J} Cn\big((Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j) \cap \mathcal{L}_{\mathcal{P}})' \cup \{\alpha\} \cup EQ_j\big) \cap \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{j\in J} Cn\big(\big(Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j)' \cap \mathcal{L}'_{\mathcal{P}}\big) \cup \{\alpha\} \cup EQ_j\big) \cap \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{j\in J} Cn((Cn(\{\alpha'\} \cup \{\beta\} \cup EQ_j)' \cup \{\alpha\} \cup EQ_j) \ \cap \\
&\qquad\qquad Cn(\mathcal{L}'_{\mathcal{P}} \cup \{\alpha\} \cup EQ_j)) \ \cap \ \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{j\in J} Cn\big((Cn\big(\{\alpha\} \cup \{\beta'\} \cup EQ'_j\big) \cup \{\alpha\} \cup EQ_j) \ \cap \ \mathcal{L}_{\mathcal{P}\cup\mathcal{P}'}\big) \ \cap \ \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{j\in J} Cn\big(Cn\big(\{\alpha\} \cup \{\beta'\} \cup EQ'_j\big) \cup \{\alpha\} \cup EQ_j\big) \ \cap \ \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{j\in J} Cn(Cn(\{\alpha\} \cup \{\beta'\} \cup EQ_j) \cup \{\alpha\} \cup EQ_j) \ \cap \ \mathcal{L}_{\mathcal{P}} \\
&= \bigcap_{j\in J} Cn(\{\alpha\} \cup \{\beta'\} \cup EQ_j) \ \cap \ \mathcal{L}_{\mathcal{P}} \\
&= \beta\dot{+}\alpha.
\end{aligned}$$

2. The proof of this part proceeds analogously to the preceding part.

3. From Part (1) above we have $(\alpha\dot{+}\beta)\dot{+}\alpha = \beta\dot{+}\alpha$.

   Since $\beta\dot{+}\alpha \vdash \alpha$ we have $\beta\dot{+}\alpha \equiv Cn\big(\alpha \wedge (\beta\dot{+}\alpha)\big)$ by propositional logic. From $(K\dot{+}4)$ we get that $Cn\big(\alpha \wedge (\beta\dot{+}\alpha)\big) \equiv \alpha\dot{+}(\beta\dot{+}\alpha)$, from which our result obtains.

■

Proof 4.10

If $K \vdash \bot$ then both parts of the theorem trivially hold.

Thus assume that $K \nvdash \bot$. Since $IC_e \cup \{\gamma\} \cup \{\alpha\} \nvdash \bot$ for every $\gamma \in IC_c$, there is a belief change extension of $(K, \{\alpha\} \cup IC_e, \overline{IC_c})$.

From Definition 4.1, we have that IC_e is true in every such extension, and every member of IC_c is consistent with every such extension. ■

A.2 Proofs of Section 5

Proof 5.1

We make use of the following Lemma.

> **Lemma A.3** *Let E_i be a belief change extension of belief change scenario $B = (K, R, C)$ with determining set of equivalences EQ_i. Then we have:*
>
> $$\vdash \left(\bigwedge_{p \equiv p' \in EQ_i} (p \equiv p') \wedge \bigwedge_{p \equiv p' \notin EQ_i} (p \equiv \neg p') \right) \supset (K' \equiv \lceil K \rceil_i).$$
>
> **Proof A.3**
>
> Let M be a model of $\bigwedge_{p \equiv p' \in EQ_i} (p \equiv p') \wedge \bigwedge_{p \equiv p' \notin EQ_i} (p \equiv \neg p')$.
>
> $\lceil K \rceil_i$ is the same as K except that for every $p \in \mathcal{P}_{\overline{EQ_i}}$, where K mentions p, $\lceil K \rceil_i$ has $\neg p$.
>
> 1. For $p \in \mathcal{P}_{EQ_i}$ we have that M assigns the same truth value to p' in K' as p in K, and so p in $\lceil K \rceil_i$.
> 2. For $p \in \mathcal{P}_{\overline{EQ_i}}$, we have that M assigns the opposite truth value to p' in K' as it does to p in K. But this means that M assigns the same truth value to p' in K' as to $\neg p$ in $\lceil K \rceil_i$.
>
> Thus M is a model of K' iff M is a model of $\lceil K \rceil_i$, from which our result follows. ■

Let $(EQ_i)_{i \in I}$ be the family of equivalences determing a belief change extension of $B = (K, \{\alpha\}, \emptyset)$. We have that

$$K \dot{+} \alpha \quad = \quad \bigcap_{i \in I} Cn(\{K'\} \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_{\mathcal{P}}.$$

As well,

$$\lceil (K, \{\alpha\}, \emptyset) \rceil \wedge \alpha \quad = \quad \bigvee_{i \in I} \lceil K \rceil_i \wedge \alpha.$$

We just need to show: For $E_i = Cn(\{K'\} \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_{\mathcal{P}}$ a belief change extension of B with determining set of equivalences EQ_i:

1. $Cn(\{K'\} \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P} \vdash \lceil K \rceil_i \wedge \alpha$ and
2. $\{\lceil K \rceil_i \wedge \alpha\} \vdash \phi$ for every $\phi \in Cn(\{K'\} \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P}$.

For each part in turn:

1. From Lemma A.3 we have

$$\vdash \left(\bigwedge_{p \equiv p' \in EQ_i} (p \equiv p') \wedge \bigwedge_{p \equiv p' \notin EQ_i} (p \equiv \neg p') \right) \supset (K' \equiv \lceil K \rceil_i).$$

Hence,

$$\{K'\} \cup EQ_i \cup \overline{EQ_i} \vdash \lceil K \rceil_i. \tag{5}$$

Since we have $\{K'\} \cup \{\alpha\} \cup EQ_i \vdash p \equiv \neg p'$ for every $(p \equiv p') \in \overline{EQ_i}$ by Theorem 4.1.1, we obtain from (5) that $\{K'\} \cup \{\alpha\} \cup EQ_i \vdash \lceil K \rceil_i$.

Hence, we get $\{K'\} \cup \{\alpha\} \cup EQ_i \vdash \lceil K \rceil_i \wedge \alpha$.

By the definition of $Cn(\cdot)$, this means that $\lceil K \rceil_i \wedge \alpha \in Cn(\{K'\} \cup \{\alpha\} \cup EQ_i)$.

Since also $\lceil K \rceil_i \wedge \alpha \in \mathcal{L}_\mathcal{P}$ we get $\lceil K \rceil_i \wedge \alpha \in Cn(\{K'\} \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P}$.

Hence $Cn(\{K'\} \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P} \vdash \lceil K \rceil_i \wedge \alpha$.

2. Assume that $\phi \in Cn(\{K'\} \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P}$.

Thus $\phi \in \mathcal{L}_\mathcal{P}$ and

$$\{K'\} \cup \{\alpha\} \cup EQ_i \vdash \phi.$$

From monotonicity of classical logic it follows that

$$\{K'\} \cup \{\alpha\} \cup EQ_i \cup \overline{EQ_i} \vdash \phi.$$

Lemma A.3 yields $\{\lceil K \rceil_i\} \cup \{\alpha\} \cup EQ_i \cup \overline{EQ_i} \vdash \phi$.

Since $\lceil K \rceil_i, \alpha, \phi \in \mathcal{L}_\mathcal{P}$ it follows that $\{\lceil K \rceil_i, \alpha\} \vdash \phi$ as required.

■

Proof 5.2

Let $(EQ_i)_{i \in I}$ determine belief change extensions of $B = (K, \emptyset, \{\alpha\})$. We have that

$$K \dot{-} \alpha \quad = \quad \bigcap_{i \in I} Cn(\{K'\} \cup EQ_i) \cap \mathcal{L}_\mathcal{P}.$$

As well,

$$\lfloor (K, \emptyset, \{\alpha\}) \rfloor \quad = \quad \bigvee_{i \in I} \bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j.$$

We just need to show that for each belief change extension of B with determining set of equivalences EQ_i:

$$Cn(\{K'\} \cup EQ_i) \cap \mathcal{L}_{\mathcal{P}} \quad \equiv \quad \bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j.$$

only-if part: We show $\{K'\} \cup EQ_i \vdash \bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j$.

Let M be a model of $\{K'\} \cup EQ_i$.

Then there is $\pi_k \in \Pi_i$ that corresponds to the assignment of truth values to members of $\mathcal{P}$ (and so $\mathcal{P}'$) in $\mathcal{P}_{\overline{EQ_i}}$; let the corresponding disjunct in $\bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j$ be $\lfloor K \rfloor_i^k$.

Since M is a model of EQ_i, for every $p \equiv p' \in EQ_i$, we obtain that M assigns the same truth values to occurrences of p' in K' as to p in $\lfloor K \rfloor_i^k$.

As well, we have chosen k so that for every $p \in \mathcal{P}_{\overline{EQ_i}}$, M assigns the opposite truth values to occurrences of p' in K' as to p in $\lfloor K \rfloor_i^k$.

Hence (using Lemma A.3) M is a model of $\lfloor K \rfloor_i^k$ and so M is a model of $\bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j$.

if part: We show that if $\{K'\} \cup EQ_i \vdash \phi$ for arbitrary $\phi \in \mathcal{L}_{\mathcal{P}}$ then $\left\{ \bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j \right\} \vdash \phi$,

or equivalently, if $\left\{ \bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j \right\} \cup \{\neg\phi\} \nvdash \bot$ then $\{K'\} \cup EQ_i \cup \{\neg\phi\} \nvdash \bot$.

So we need to find a model M, over the language $\mathcal{L}_{\mathcal{P} \cup \mathcal{P}'}$, of $\left\{ \bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j \right\} \cup \{\neg\phi\}$ such that M is also a model of $\{K'\} \cup EQ_i \cup \{\neg\phi\}$.

Let $M^{\mathcal{P}}$ be a model over $\mathcal{L}_{\mathcal{P}}$ of $\left\{ \bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j \right\} \cup \{\neg\phi\}$.

For $p \in \mathcal{P}_{\overline{EQ_i}}$, $M^{\mathcal{P}}$ coincides with a specific mapping, $\pi_k \in \Pi_i$. As well, $M^{\mathcal{P}}$ satisfies a specific disjunct $\lfloor K \rfloor_i^k$ of $\bigvee_{\pi_j \in \Pi_i} \lfloor K \rfloor_i^j$.

We extend $M^{\mathcal{P}}$ to a model M over $\mathcal{L}_{\mathcal{P} \cup \mathcal{P}'}$ as follows.

1. M is the same as $M^{\mathcal{P}}$ for atoms in $\mathcal{P}$.
2. For $p \equiv p' \in EQ_i$, M assigns the same value to p' as $M^{\mathcal{P}}$ does to p.
3. The remaining atoms $p' \in \mathcal{P}'$ (and so for $p \equiv p' \in \overline{EQ_i}$) are assigned according to π_k's assignment to atoms of $\mathcal{P}$.

Thus from 1. we get that $\neg\phi$ is satisfied; from 2. we get that EQ_i is satisfied; and from 3. we get that K' is satisfied.

■

Proof 5.3

We make use of the following Lemmas.

Lemma A.4 *Let EQ be a set of equivalences determining a consistent belief change extension of belief change scenario $B = (K, R, C)$.*

Then $\{p \equiv p' \mid p \in \mathcal{P}(K) \,\Delta\, \mathcal{P}(R \cup C)\} \subseteq EQ$.

(So if p is mentioned in K, but not R or C, or else in R or C, but not K, then $p \equiv p' \in EQ$ for any EQ determining a belief change extension of $B = (K, R, C)$.

Proof A.4

Assume otherwise. So there is a belief change extension of belief change scenario $B = (K, R, C)$ where for corresponding set of equivalences EQ we have

1. $\exists p \in \mathcal{P}$ where $p \in \mathcal{P}(K)$, $p \notin \mathcal{P}(R \cup C)$ and $p \equiv p' \notin EQ$ or
2. $\exists p \in \mathcal{P}$ where $p \notin \mathcal{P}(K)$, $p \in \mathcal{P}(R \cup C)$ and $p \equiv p' \notin EQ$.

For the first case, and for p as above, we have from Theorem 4.1.1 that for some $\phi \in C \cup \{\bot\}$ that $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash \neg(p \equiv p')$ or $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash (p \vee p') \wedge (\neg p \vee \neg p')$.

So:

$$K' \cup R \cup \{\neg\phi\} \cup EQ \;\vdash\; p \vee p' \tag{6}$$
$$K' \cup R \cup \{\neg\phi\} \cup EQ \;\vdash\; \neg p \vee \neg p' \tag{7}$$

We have that $p \notin \mathcal{P}(R \cup \{\neg\phi\})$ by assumption, and clearly $p \notin \mathcal{P}(EQ)$, and $p \notin \mathcal{P}(K')$. That is, p doesn't appear on the left hand side of $\vdash$ in (6) and (7).

So from (6) we must have

$$K' \cup R \cup \{\neg\phi\} \cup EQ \vdash p'. \tag{8}$$

(If this isn't the case and $K' \cup R \cup \{\neg\phi\} \cup EQ \nvdash p'$ then there is a model M of $K' \cup R \cup \{\neg\phi\} \cup EQ$ that isn't a model of p'. Let M_1 be the same as M but assigning $false$ to p. Then M_1 is a model of $K' \cup R \cup \{\neg\phi\} \cup EQ$ but not of $p \vee p'$, contradicting (6).)

Analogously, from (7) we derive

$$K' \cup R \cup \{\neg\phi\} \cup EQ \vdash \neg p'. \tag{9}$$

But (8) + (9) gives $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash \bot$, contradicting the fact that EQ determines a belief change extension.

For the second case, we derive, as previously,

$$\begin{aligned} K' \cup R \cup \{\neg\phi\} \cup EQ &\vdash p \vee p' \\ K' \cup R \cup \{\neg\phi\} \cup EQ &\vdash \neg p \vee \neg p' \end{aligned}$$

Since $p \notin \mathcal{P}(K)$ by assumption, so $p' \notin \mathcal{P}(K')$. As well, clearly $p' \notin \mathcal{P}(R \cup \{\neg\phi\})$ and $p' \notin \mathcal{P}(EQ)$. Analogous to the first case, we obtain the contradiction

$$\begin{aligned} K' \cup R \cup \{\neg\phi\} \cup EQ &\vdash \neg p \\ K' \cup R \cup \{\neg\phi\} \cup EQ &\vdash p. \end{aligned}$$

■

Lemma A.5 *Let $B = (K, R, C)$ be a belief change scenario and let $\mathcal{Q} = \mathcal{P}(K) \cap \mathcal{P}(R \cup C)$.*

For $EQ \subseteq \{p \equiv p' \mid p \in \mathcal{P}\}$ a set of equivalences determining some belief change extension of B we have that, for every $\phi \in C \cup \{\bot\}$:

$$Cn(K' \cup R \cup \{\neg\phi\} \cup EQ) \cap \mathcal{L}_{\mathcal{P}} \;=\; Cn\big(K'[\mathcal{Q}] \cup R \cup \{\neg\phi\} \cup EQ^{\mathcal{Q}}\big) \cap \mathcal{L}_{\mathcal{P}}$$

where $EQ^{\mathcal{Q}} = EQ \setminus \{p \equiv p' \mid p \notin \mathcal{Q}\}$.

Thus there is a 1–1 correspondence between sets EQ and $EQ^{\mathcal{Q}}$ determining belief change extensions of belief change scenario B.

Proof A.5

Only-if part: We show that any model of $K' \cup R \cup \{\neg\phi\} \cup EQ$ is also a model of $K'[\mathcal{Q}] \cup R \cup \{\neg\phi\} \cup EQ^{\mathcal{Q}}$.

Let M be a model of $K' \cup R \cup \{\neg\phi\} \cup EQ$. From substitution of equivalent formulas (here in EQ) we get that $K' \cup EQ \vdash K'[\mathcal{Q}]$.

Thus since M is a model of $K' \cup EQ$ it is of $K'[\mathcal{Q}]$.

Since $EQ^{\mathcal{Q}} \subseteq EQ$ (Lemma A.4), and M is a model of EQ, it is also of $EQ^{\mathcal{Q}}$. Thus M is a model of $K'[\mathcal{Q}] \cup R \cup \{\neg\phi\} \cup EQ^{\mathcal{Q}}$.

If part: We show that for arbitrary $\delta \in \mathcal{L}_{\mathcal{P}}$, any proof of $K'[\mathcal{Q}] \cup R \cup \{\neg\phi\} \cup EQ^{\mathcal{Q}} \vdash \delta$ can be transformed into a proof of $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash \delta$.

Let $\psi_1, \ldots, \psi_n = \delta$ be a proof of δ from $K'[\mathcal{Q}] \cup R \cup \{\neg\phi\} \cup EQ^{\mathcal{Q}}$.

We construct a proof of δ from $K' \cup R \cup \{\neg\phi\} \cup EQ$ as follows.

For ψ_j, $1 \leq j \leq n$ we have the following cases.

1. $\vdash \psi_j$. We leave ψ_j unchanged.

2. $\psi_j \in K'[\mathcal{Q}]$. It follows easily that $K' \cup EQ \vdash \psi_j$. We replace ψ_j by a proof (sequence of formulas) of ψ_j from $K' \cup EQ$.
3. $\psi_j = R$ or $\psi_j = \{\neg\phi\}$. We leave ψ_j unchanged.
4. $\psi_j \in EQ^{\mathcal{Q}}$. We leave ψ_j unchanged.
5. ψ_j results from ψ_k, ψ_l, $1 \leq k, l < j$ by an application of modus ponens. Since, by an induction hypotheses, we have ψ_k, ψ_l are logical consequences of $K' \cup R \cup \{\neg\phi\} \cup EQ$ and ψ_l is $\psi_k \supset \psi_j$, we obtain $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash \psi_j$ by modus ponens.

Hence we obtain a sequence of formulas where each formula is

1. a tautology,
2. a premiss drawn from the set $K' \cup R \cup \{\neg\phi\} \cup EQ$, or
3. obtained from previous formulas in the sequence by an application of modus ponens.

Hence we have shown that $K' \cup R \cup \{\neg\phi\} \cup EQ \vdash \delta$.

■

Let $B = (K, \{\alpha\}, \emptyset)$ be a belief change scenario.

1. For $\dot{+}_c$ we have:

   Let E_i be a belief change extension of B such that $K\dot{+}_c\alpha = E_i$ for selection function c. Then we have for EQ_i determining E_i that:

   $$K' \cup \{\alpha\} \cup EQ_i \nvdash \bot \qquad \text{and}$$

   $$E_i = Cn(K' \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_{\mathcal{P}}.$$

   From Lemma A.5 we obtain that

   $$Cn(K' \cup \{\alpha\} \cup EQ_i) \cap \mathcal{L}_{\mathcal{P}} \;=\; Cn\big(K'[\mathcal{Q}] \cup \{\alpha\} \cup EQ_i^{\mathcal{Q}}\big) \cap \mathcal{L}_{\mathcal{P}}.$$

   Hence $K'[\mathcal{Q}] \cup \{\alpha\} \cup EQ_i^{\mathcal{Q}} \nvdash \bot$.

   As well for $p \equiv p' \notin EQ_i$ we obtain $K'[\mathcal{Q}] \cup \{\alpha\} \cup EQ_i^{\mathcal{Q}} \cup \{p \equiv p'\} \vdash \bot$, again from Lemma A.5.

   Thus $E_i = K\dot{+}_c^{\mathcal{Q}}\alpha$ is a choice revision for selection function c with respect to $\mathcal{Q}$.

2. For $\dot{+}$, the theorem follows by noting that for belief change scenario B, there is a 1-1 correspondence between belief change extensions determined by a set of equivalences EQ_i and the corresponding set $EQ_i^{\mathcal{Q}}$.

3. Proofs for $\dot{-}_c$ and $\dot{-}_c^v$ follow analogously to those for revision.

■

Proof 5.4

From Theorem 5.1 we have that $K\dot{+}\alpha \equiv \bigvee_{i\in I} \lceil K \rceil_i \wedge \alpha$.

Hence we just need to show that $\bigvee_{i\in I} \lceil K \rceil_i \wedge \alpha. \equiv . \bigvee_{i\in I} \lceil K \rceil_i^{\mathcal{Q}} \wedge \alpha$.

We have, for any EQ_i determining a belief change extension of B, that $p \in \mathcal{P}(K)\,\Delta\,\mathcal{P}(\alpha)$ implies that $p \equiv p' \in EQ_i$, or $p \equiv p' \in \overline{EQ_i}$ implies that $p \in \mathcal{P}(K) \cap \mathcal{P}(\alpha)$.

Thus EQ_i determines a belief change extension of B (via Definition 4.1) iff EQ_i determines a belief change extension of B with respect to $\mathcal{Q}$ (via Definition 5.3).

From this it follows that (informally stated) Definitions 5.1 and 5.5 identify precisely the same formulas, from which our result follows.

■

Proof 5.5

The proof relies upon the following lemma which follows easily from the results proven in [Rei80].

Lemma A.6 *Let W and E be sets of formulas and let D be a set of default rules of the form $\frac{:\beta,\phi_1,\ldots,\phi_n}{\beta}$ where $\beta, \phi_1, \ldots, \phi_n$ are formulas.*

Then, we have that E is an extension of (D, W) iff $E = Cn\Big(W \cup \{\beta \mid \frac{:\beta,\phi_1,\ldots,\phi_n}{\beta} \in D'\}\Big)$ for some maximal subset $D' \subseteq D$ such that for every $\frac{:\beta,\phi_1,\ldots,\phi_n}{\beta} \in D'$ we have that $\neg\beta \notin E$ and $\neg\phi_i \notin E$ for $i = 1..n$.

Moreover, [Rei80] tells us that E is consistent iff W is consistent.

Let $B = (K, R, C)$ be a belief change scenario. Define

$$\Delta_B = \left(\left\{ \frac{:\ p \equiv p', \neg\phi_1, \ldots, \neg\phi_n}{p \equiv p'} \;\middle|\; p \in \mathcal{P} \right\}, K' \cup R \right).$$

Assume B is an inconsistent belief change scenario, that is, $K' \cup R$ is inconsistent. Then, by Definition 4.1, we have that $\mathcal{L}_{\mathcal{P}}$ is the sole (inconsistent) belief change extension of B. According to [Rei80], the inconsistency of $K' \cup R$ implies that Δ_B has a single (inconsistent) extension $\mathcal{L}_{\mathcal{P}\cup\mathcal{P}'}$.

For the remainder, assume that $K' \cup R$ is consistent.

only-if part: Let E be an extension of Δ_B.

According to Lemma A.6, we have that

$$E = Cn\Big(K' \cup R \cup \Big\{(p \equiv p') \mid \tfrac{:(p\equiv p'),\phi_1,\ldots,\phi_n}{(p\equiv p')} \in D'\Big\}\Big)$$

for some maximal subset $D' \subseteq \left\{ \frac{:p\equiv p',\phi_1,\ldots,\phi_n}{p\equiv p'} \;\middle|\; p \in \mathcal{P} \right\}$ such that $\neg(p \equiv p') \notin E$ and $\neg\phi_i \notin E$ for $i = 1..n$.

We show that $EQ = \{(p \equiv p') \mid \frac{:(p\equiv p'),\phi_1,\ldots,\phi_n}{(p\equiv p')} \in D'\}$ determines a belief change extension F of B such that $F = E \cap \mathcal{L}_{\mathcal{P}}$.

In fact, $E = Cn(K' \cup R \cup EQ)$. By the theory of default logic, we get that $E = Cn(K' \cup R \cup EQ)$ is consistent, due to the consistency of $K' \cup R$. That is, $\bot \notin Cn(K' \cup R \cup EQ)$.

Moreover, we get that EQ is maximal in satisfying

$$Cn(K' \cup R \cup EQ) \cap C = \emptyset \, .$$

As a consequence, EQ determines the belief change extension F of B.

if part: Let F be a belief change extension of B determined by EQ. Define $E = Cn(K' \cup R \cup EQ)$. Clearly, we have $F = E \cap \mathcal{L}_{\mathcal{P}}$.

By definition, EQ is a maximal set of equivalences satisfying

$$Cn(K' \cup R \cup EQ) \cap \{\phi_1, \ldots, \phi_n, \bot\} = \emptyset.$$

That is, $\neg\phi_i \notin E$ for $i = 1..n$. Clearly, we also have $\neg(p \equiv p') \notin E$ for all $(p \equiv p') \in EQ$.

Given that EQ is also maximal with respect to the latter requirements, it induces a maximal subset $D' \subseteq \left\{ \frac{: p \equiv p', \phi_1, \ldots, \phi_n}{p \equiv p'} \;\middle|\; (p \equiv p') \in EQ \right\}$ such that

$$E = Cn\Big(K' \cup R \cup \Big\{(p \equiv p') \mid \tfrac{: (p \equiv p'), \phi_1, \ldots, \phi_n}{(p \equiv p')} \in D'\Big\}\Big) \, .$$

According to Lemma A.6, E is an extension of Δ_B. ■

Proof 5.6 See [DSTW01]. ■

References

[AGM85] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[AM82] C.E. Alchourrón and D. Makinson. On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.

[BDP01] S. Benferhat, D. Dubois, and H. Prade. A computational model for belief change and fusing ordered belief bases. In M.-A. Williams and H. Rott, editors, *Frontiers in Belief Revision*, volume 22 of *Applied Logic Series*, pages 109–134. Kluwer Academic Publishers, 2001.

[BG93] C. Boutilier and M. Goldszmidt. Revision by conditional beliefs. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 649–654. Morgan Kaufmann, Washington, D.C., July 1993.

[Bor85] A. Borgida. Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems*, 10, 1985.

[Bou94] C. Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence*, 68(1):33–85, 1994.

[BS98] P. Besnard and T. Schaub. Signed systems for paraconsistent reasoning. *Journal of Automated Reasoning*, 20(1-2):191–213, 1998.

[CW94] T. Chou and M. Winslett. A model-based belief revision system. *Journal of Automated Reasoning*, 12(2):157–208, 1994.

[Dal88] 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.

[DHS02] J. Delgrande, A. Hunter, and T. Schaub. Coba: A consistency-based belief revision system. In S. Greco and N. Leone, editors, *Proceedings of the European Workshop on Logics in Artificial Intelligence (JELIA 2002)*, Lecture Notes in Artificial Intelligence. Springer Verlag, 2002.

[dK86] J. de Kleer. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.

[dK89] J. de Kleer and K. Konolige. Eliminating the fixed predicates from a circumscription. *Artificial Intelligence*, 39(3):391–398, July 1989.

[DP97] A. Darwiche and J. Pearl. On the logic of iterated revision. *Artificial Intelligence*, 89:1–29, 1997.

[DS00] J.P. Delgrande and T. Schaub. A consistency-based model for belief change: Preliminary report. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 392–398, Austin, TX, 2000.

[DS02] J.P. Delgrande and T. Schaub. A consistency-based approach for belief change: Full report. Technical report, School of Computing Science, Simon Fraser University and Institut für Informatik Universität Potsdam, 2002.

[DSTW01] 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.

[dV93] A. del Val. Syntactic characterizations of belief change operators. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 540–545, Chambéry, France, 1993.

[EETW00] U. Egly, T. Eiter, H. Tompits, and S. Woltran. Solving advanced reasoning tasks using quantified Boolean formulas. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 417–422, Austin, TX, 2000.

[EG92] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57(2-3):227–270, 1992.

[For89] K.D. Forbus. Introducing actions into qualitative simulation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1273–1278, 1989.

[Fuh88] A. Fuhrmann. *Relevant Logics, Modal Logics, and Theory Change*. PhD thesis, Australian National University, Australia, 1988.

[FUV83] R. Fagin, J. Ullman, and M. Vardi. On the semantics of updates in databases. preliminary report. In *Proceedings of the Second ACM Symposium on Principles of Database Systems SIGACT-SIGMOD*, pages 352–365, 1983.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.

[GLMT02] E. Giunchiglia, J. Lee, V. Lifschitz N. McCain, and H. Turner. Nonmonotonic causal theories. unpublished draft, 2002.

[Gro88] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[Han99] Sven Ove Hansson. *A Textbook of Belief Dynamics*. Applied Logic Series. Kluwer Academic Publishers, 1999.

[KM91] H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.

[KM92] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In P. Gärdenfors, editor, *Belief Revision*, pages 183–203. Cambridge University Press, 1992.

[Kow78] R. Kowalski. Logic for data description. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 77–103. Plenum Press, 1978.

[Leh95] D. Lehmann. Belief revision, revised. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1534–1540, Montréal, 1995. Morgan Kaufmann Publishers.

[Lib99] P. Liberatore. BReLS: A system for revising, updating, and merging knowledge bases. In *The Third IJCAI Workshop on Nonmonotonic Reasoning, Action and Change*, pages 41–48, Stockholm, Sweden, August 1999.

[Lin00] F. Lin. On strongest necessary and weakest sufficient conditions. In A. Cohn, F. Giunchiglia, and B. Selman, editors, *Proceedings of the Seventh International Conference on the Principles of Knowledge Representation and Reasoning*, pages 167–175. Morgan Kaufmann Publishers, 2000.

[LR94] F. Lin and R. Reiter. Forget it! In *AAAI Fall Symposium on Relevance*, New Orleans, November 1994.

[LS97] P. Liberatore and M. Schaerf. Reducing belief revision to circumscription (and vice versa). *Artificial Intelligence*, 93(1–2):261–296, 1997.

[LS00] P. Liberatore and M. Schaerf. Brels: A system for the integration of knowledge bases. In A. Cohn, F. Giunchiglia, and B. Selman, editors, *Proceedings of the Seventh International Conference on the Principles of Knowledge Representation and Reasoning*, pages 145–152. Morgan Kaufmann Publishers, 2000.

[Mak85] D. Makinson. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthese*, 62:347–363, 1985.

[Neb92] B. Nebel. Syntax based approaches to belief revision. In P. Gärdenfors, editor, *Belief Revision*, pages 52–88. Cambridge University Press, 1992.

[NFPS96] A.C. Nayak, N.Y. Foo, M. Pagnucco, and A. Sattar. Changing conditional belief unconditionally. In *Proc. of the Sixth Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 119–136, De Zeeuwse Stromen, The Netherlands, 1996.

[Pap01] O. Papini. Iterated revision operations stemming from the history of an agent's observations. In M.-A. Williams and H. Rott, editors, *Frontiers in Belief Revision*, volume 22 of *Applied Logic Series*, pages 279–301. Kluwer Academic Publishers, 2001.

[Par99] Rohit Parikh. Beliefs, belief revision, and splitting languages. In L.S. Moss, J. Ginzburg, and M. de Rijke, editors, *Logic, Language and Computation, Vol II*, pages 266–278. CSLI Publications, 1999.

[Poo88] D.L. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36(1):27–48, 1988.

[RdK87] R. Reiter and J. de Kleer. Foundations for assumption-based truth maintenance systems: Preliminary report. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 183–188. Morgan Kaufmann Publishers, 1987.

[Rei80] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[Rei84] 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.

[Rei87a] R. Reiter. On integrity constraints. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 97–111, Pacific Grove, CA, 1987.

[Rei87b] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–96, 1987.

[Rot01] H. Rott. *Change, Choice and Inference - A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford: Clarendon Press, 2001.

[Sat88] K. Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pages 455–462, Tokyo, 1988.

[SK87] 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.

[Som94] Léa Sombé, editor. *Revision and Updating in Knowledge Bases*. John Wiley and Sons, Inc., 1994.

[Spo88] 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.

[Web86] A. Weber. Updating propositional formulas. *Proc. First Conference on Expert Database Systems*, pages 487–500, 1986.

[Wil94] M.-A. Williams. Transmutations of knowledge systems. In J. Doyle, P. Torasso, and E. Sandewall, editors, *Proceedings of the Fourth International Conference on the Principles of Knowledge Representation and Reasoning*, pages 619 – 629, Bonn, Germany, May 1994.

[Wil95] M.-A. Williams. Iterated theory base change: A computational model. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1541–1547, Montréal, 1995.

[Win88] M. Winslett. Reasoning about action using a possible models approach. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 89–93, St. Paul, Minnesota, 1988.