Journal of Logic, Language, and Information **9:** 353–377, 2000.

Infobase Change: A First Approximation

THOMAS ANDREAS MEYER*
Department of Computer Science, University of Pretoria, Pretoria 0002, South Africa
E-mail: tmeyer@cs.up.ac.za

WILLEM ADRIAN LABUSCHAGNE
Department of Computer Science and Information Systems, University of South Africa, Box 392, Pretoria 0003, South Africa
E-mail: labuswa@alpha.unisa.ac.za

JOHANNES HEIDEMA
Department of Mathematics, Applied Mathematics and Astronomy, University of South Africa, Box 392, Pretoria 0003, South Africa
E-mail: heidej@alpha.unisa.ac.za

(Received 1 December 1997; in final form 3 November 1998)

Abstract. Generalisations of theory change involving operations on arbitrary sets of wffs instead of on belief sets (i.e., sets closed under a consequence relation), have become known as base change. In one view, a base should be thought of as providing more structure to its generated belief set, which means that it can be employed to determine the theory contraction operation associated with a base contraction operation. In this paper we follow such an approach as the first step in defining *infobase change*. We think of an infobase as a finite set of wffs consisting of independently obtained bits of information. Taking AGM theory change (Alchourrón et al., 1985) as the general framework, we present a method that uses the structure of an infobase B to obtain an AGM theory contraction operation for contracting the belief set $Cn(B)$. Both the infobase and the obtained theory contraction operation then play a role in constructing a unique infobase contraction operation. Infobase revision is defined in terms of an analogue of the Levi Identity, and it is shown that the associated theory revision operation satisfies the AGM postulates for revision. Because every infobase is associated with a unique infobase contraction and revision operation, the method also allows for iterated base change.

Key words: Base change, base contraction, base revision, belief contraction, belief revision, iterated base change, theory change

1. Introduction

One of the dominant approaches to theory change, the AGM approach (Alchourrón et al., 1985), operates on belief sets – sets of wffs of a logic language closed under a consequence operation. It is generally accepted that belief sets do not have a

* Financial assistance by the Centre for Science Development for this research is hereby acknowledged. Opinions expressed, or conclusions reached in this publication, are those of the author and should not necessarily be ascribed to the Centre for Science Development.

rich enough structure to serve as appropriate models for epistemic states (Hansson, 1992b; Gärdenfors, 1988: 67), and AGM theory change is therefore regarded as an elegant idealisation of a more general theory of belief change.* The proposal to replace the contraction of belief sets with the contraction of arbitrary sets of wffs (or bases) has come to be known as base contraction. Base contraction is usually viewed as a generalisation of theory contraction in which the contraction of belief sets is a special case. Accordingly, the methods for constructing base contraction operations are appropriate generalisations of methods for constructing AGM theory contraction operations (Hansson, 1989, 1992a, 1993; Fuhrmann, 1991; Nayak, 1994). A different (though not completely unrelated) view is that a base should be thought of as providing more structure to its associated belief set. Let us define the theory contraction operation – *associated* with a base contraction operation $\sim$ as: $Cn(B) - \phi = Cn(B \sim \phi)$. The added structure of the base can be used, in one way or another, to pick an appropriate associated theory contraction operation, from which the base contraction operation can then be constructed. This is the view encountered in Nebel's description of base contraction (Nebel, 1989, 1990, 1991, 1992).

In this paper we follow the latter approach as a first step in defining *infobase change*. We think of an infobase as a finite set of wffs, consisting of independently obtained bits of information. Taking AGM theory change as the general framework in which to operate, and viewing contraction as more primitive than revision, we present a method that uses the structure of an infobase B to determine which AGM theory contraction operation to associate with the infobase contraction operation to be constructed. This places us in a position to determine which wffs in B to retain (and also which wffs cannot be retained) during a contraction of B. Instead of simply discarding the wffs of B that cannot be retained, as most approaches to base contraction do, we rather replace them with appropriately weakened wffs. Infobase revision is defined in terms of infobase contraction by means of an infobase change analogue of the Levi Identity. We give results to indicate that our method for constructing infobase change operations is satisfactory as a first approximation, and we compare infobase change with the related approaches of Nebel (1989) and Nayak (1994). To conclude, we provide suggestions for further developments.

1.1. PRELIMINARIES

For the rest of this paper L denotes any logic language, closed under the usual propositional connectives, and containing the symbols $\top$ and $\bot$. For every finite $C, D \subseteq L$ we write $C \Diamond D$ as an abbreviation for $\{\gamma \Diamond \delta \mid \gamma \in C \text{ and } \delta \in D\}$ where $\Diamond \in \{\vee, \wedge\}$, $\neg C$ as an abbreviation for $\{\neg\gamma \mid \gamma \in C\}$, $\bigwedge C$ as an abbreviation for the conjunction of all elements in C, with $\bigwedge \emptyset = \top$, and $\bigvee C$ as an abbreviation for the disjunction of all elements in C, with $\bigvee \emptyset = \bot$. We assume

* Although the original AGM postulates are not exclusively concerned with belief sets, the major results in Alchourrón et al. (1985) only hold for belief sets.

L to have a two-valued model-theoretic semantics defining truth and falsity. The set of interpretations of L is denoted by U. We use $\Vdash$ for the relation from U to L denoting satisfaction and we assume that $\Vdash$ behaves classically with respect to the propositional connectives. We use $\top$ and $\bot$ as canonical representatives for the logically valid and logically invalid wffs respectively. For concreteness the reader may think of the logic under consideration as a (possibly infinitely generated) propositional logic. For every $X \subseteq L$, we denote the set of *models* of X by $M(X)$, and for $\alpha \in L$ we write $M(\alpha)$ instead of $M(\{\alpha\})$. Classical entailment (from $\wp L$ to L) is denoted by $\vDash$, and for $\alpha, \beta \in L$ we write $\alpha \vDash \beta$ instead of $\{\alpha\} \vDash \beta$. We also require $\vDash$ to satisfy *compactness*.* Closure under entailment is denoted by Cn. A *theory* or a *belief set* is a set $K \subseteq L$ closed under entailment. For every $V \subseteq U$, we let $Th(V)$ denote the *theory determined by* V, and for $x \in U$ we write $Th(x)$ instead of $Th(\{x\})$. A set $X \subseteq L$ *axiomatises* a set of interpretations V iff $Cn(X) = Th(V)$. For a set $X \subseteq L$, the *expansion* of X by a wff $\alpha \in L$ is defined as $X + \alpha = Cn(X \cup \{\alpha\})$.

Our examples are phrased in propositional languages, containing the usual propositional connectives, that are generated by at most three atoms. We use the letters p, q and r to denote these atoms, and interpretations of the languages will be represented by appropriate sequences of 0s and 1s, 0 representing falsity and 1 representing truth. The convention is that the first digit in the sequence represents the truth value of p, the second the truth value of q and the third the truth value of r.

2. AGM Theory Change

AGM theory contraction and revision can be described in terms of sets of postulates. Since all the AGM postulates deal with fixed belief sets, we assume a fixed belief set K and define K-contraction and K-revision functions as functions from L to the set of belief sets. Where there is no ambiguity, we shall drop the references to K. The eight postulates for AGM revision are given below.

$(K * 1)$ $K * \phi = Cn(K * \phi)$,

$(K * 2)$ $\phi \in K * \phi$,

$(K * 3)$ $K * \phi \subseteq K + \phi$,

$(K * 4)$ If $\neg\phi \notin K$, then $K + \phi \subseteq K * \phi$,

$(K * 5)$ $K * \phi = L$ iff $\vDash \neg\phi$,

$(K * 6)$ If $\phi \equiv \psi$ then $K * \phi = K * \psi$,

$(K * 7)$ $K * (\phi \wedge \psi) \subseteq (K * \phi) + \psi$,

* That is, for every $X \subseteq L$ and every $\alpha \in L$, $X \vDash \alpha$ iff $X_F \vDash \alpha$ for some finite subset X_F of X.

$(K*8)$ If $\neg\psi \notin K*\phi$, then $(K*\phi)+\psi \subseteq K*(\phi \wedge \psi)$.

The AGM contraction postulates follow a similar pattern.

$(K-1)$ $K-\phi = Cn(K-\phi)$,

$(K-2)$ $K-\phi \subseteq K$,

$(K-3)$ If $\phi \notin K$ then $K-\phi = K$,

$(K-4)$ If $\not\models \phi$ then $\phi \notin K-\phi$,

$(K-5)$ If $\phi \in K$ then $(K-\phi)+\phi = K$,

$(K-6)$ If $\phi \equiv \psi$ then $K-\phi = K-\psi$,

$(K-7)$ $(K-\phi) \cap (K-\psi) \subseteq K-(\phi \wedge \psi)$,

$(K-8)$ If $\psi \notin K-(\phi \wedge \psi)$ then $K-(\phi \wedge \psi) \subseteq K-\psi$.

The following two identities can be used to define revision and contraction in terms of one another.

(Harper Identity) $K-\phi = K \cap (K*\neg\phi)$

(Levi Identity) $K*\phi = (K-\neg\phi)+\phi$

From results in Grove (1988), Katsuno and Mendelzon (1991) and Boutilier (1994), AGM theory change can be characterised by a set of preorders on U.

DEFINITION 1. Let $\preceq$ be any preorder (i.e., a reflexive transitive relation) on U.

1. $x \in V \subseteq U$ is $\preceq$-*minimal* in V iff for every $y \in V$, $y \not\prec x$.
2. For a $V \subseteq U$, $\preceq$ is V-*smooth* iff for every $y \in V$ there is an $x \preceq y$ that is $\preceq$-minimal in V.
3. $\preceq$ is smooth iff $\preceq$ is $M(\phi)$-smooth for every ϕ. We denote the set of $\preceq$-minimal elements of $M(\phi)$ by $\text{Min}_{\preceq}(\phi)$.
4. Given an arbitrary set $X \subseteq L$, a preorder $\preceq$ on U is X-*faithful* iff $\preceq$ is smooth, $x \prec y$ for every $x \in M(X)$ and $y \notin M(X)$, and $x \not\prec y$ for every $x, y \in M(X)$.

The idea is to consider preorders in which the models of X, being the minimal, or "best" interpretations, are strictly below all other interpretations.

DEFINITION 2.

1. A revision function $*$ is *obtained from* a K-faithful preorder $\preceq$ iff $K*\phi = Th(\text{Min}_{\preceq}(\phi))$.
2. A contraction function $-$ is *obtained from* a K-faithful preorder $\preceq$ iff $K-\phi = Th(M(K) \cup \text{Min}_{\preceq}(\neg\phi))$.

A reinterpretation of Grove's "systems of spheres" (1988) as K-faithful *total* preorders yields a semantic characterisation of AGM theory change, in the same spirit as that of Katsuno and Mendelzon (1991), and Boutilier (1994).⋆ We state these results without proof in the theorem below, and use it throughout the rest of this paper without explicit references to it.

THEOREM 1.

1. *A revision function satisfies* $(K * 1)$ *to* $(K * 8)$ *iff it is obtained from some* K*-faithful total preorder.*
2. *A contraction function satisfies* $(K - 1)$ *to* $(K - 8)$ *iff it is obtained from some* K*-faithful total preorder.*

3. Infobase Change

In this section we present an approach to base change that relies heavily on the structure of the base. Intuitively, we regard every wff in a base B as an explicit piece of information, obtained "independently" from the other wffs in B. Such a base will be referred to as an *infobase*. We assume infobases to be finite. Wffs in an infobase B should be seen as the primary sources of information, with the wffs in $Cn(B) \setminus B$ as secondary or derived bits of information, owing their status as beliefs only to the fact that they are entailed by one or more wffs in B. Our approach takes the syntactic form of infobases into account, since two different infobases corresponding to the same belief set may behave differently, but the syntactic form of the wffs in an infobase is irrelevant. We take infobase contraction as more primitive than infobase revision, preferring to define infobase revision in terms of infobase contraction by means of an infobase change analogue of the Levi Identity. Formally, we assume a fixed infobase B, and define infobase B-contraction and infobase B-revision functions as functions from L to $\wp L$. Where there is no room for ambiguity, we shall omit references to B.

3.1. Infobase Contraction

To construct an infobase B-contraction operation, we first use the wffs in B to induce a B-faithful total preorder. The theory $Cn(B)$-contraction function obtained from the B-faithful total preorder is taken to be the theory contraction function associated with the infobase contraction function that we aim to construct. Using the intuition associated with an infobase, we obtain the appropriate B-faithful total preorder by ordering the interpretations according to the number of non-equivalent wffs of B they satisfy – the more they satisfy, the "better" they are deemed to be, and the lower down in the total preorder they will be.

⋆ In fact, our account is slightly more general than that of Grove. We allow for the existence of elementarily equivalent interpretations whereas he does not. But to obtain this added generality requires only a few small modifications.

DEFINITION 3. For every $x \in U$ we define x_B, the *B-number of x*, as the number of logically non-equivalent wffs in B that are true in x. The *B-induced* faithful total preorder $\preceq$ on U is defined as: $x \preceq y$ iff $x_B \geq y_B$. The *B-induced contraction function* $-$ is defined as $Cn(B) - \phi = Th(M(B) \cup \text{Min}_{\preceq}(\neg\phi))$.

It is easy to verify that $\preceq$ is a B-faithful total preorder and that $-$ therefore satisfies the AGM contraction postulates.*

Taking the B-induced contraction function as the theory contraction function associated with the infobase contraction function allows us to determine which wffs in B should be retained, and which cannot be retained, after a contraction of B.

DEFINITION 4. The set of *ϕ-discarded* wffs (of B) is defined as $B^{-\phi} = \{\psi \in B \mid \psi \notin Cn(B) - \phi\}$, where $-$ is the B-induced contraction function. We refer to $B \setminus B^{-\phi}$ as the *ϕ-retained wffs* (of B).

Clearly, the ϕ-retained wffs are precisely the wffs in B that should be retained when contracting B by ϕ. Unlike most approaches to base contraction, we do not simply expunge the ϕ-discarded wffs, but instead opt to replace them with appropriately weakened wffs. The strategy is to retain as much of the information contained in a wff as possible, even if not all the information in the wff can be retained. This is in line with the intuition that infobases consist of independently obtained wffs. Of course, these weakened wffs cannot be chosen in an arbitrary fashion. Since the B-induced contraction function $-$ has already been identified as the theory contraction function to be associated with the infobase contraction function, the weakened wffs, together with the ϕ-retained wffs, have to generate the belief set $Cn(B) - \phi$.

A first attempt at weakening might be to add the minimal models of $\neg\phi$ to the models of every ϕ-discarded wff, and if each of the resulting sets of interpretations is axiomatisable, to let the corresponding wffs be the appropriate weakened versions. However, as the next example shows, this proposal does not quite do justice to the intuition concerning the independence of wffs in B.

Example 1. Let $B = \{p, q, r\}$. Figure 1 gives a graphical representation of the B-induced faithful total preorder $\preceq$. Because p, q and r each represents independently obtained information, a contraction with $p \wedge q$ should have no effect on r. That is, when contracting B by $p \wedge q$, the resulting infobase should contain weakened versions of p and q, and should contain r itself. For the same

* The definition of B-induced contraction functions flies somewhat in the face of the idea that the wffs in an infobase are independently obtained, since the B-induced faithful total preorders do not distinguish between independently obtained wffs that are logically equivalent. But to be able to draw such a distinction it is not sufficient to represent an infobase as a *set* of wffs; a tuple or list is more appropriate. This is because two or more instances of the same wff might also be thought of as independently obtained. We shall, for the rest of this paper, ignore this desirable refinement of infobase change and stick to the representation of infobases as sets of wffs.

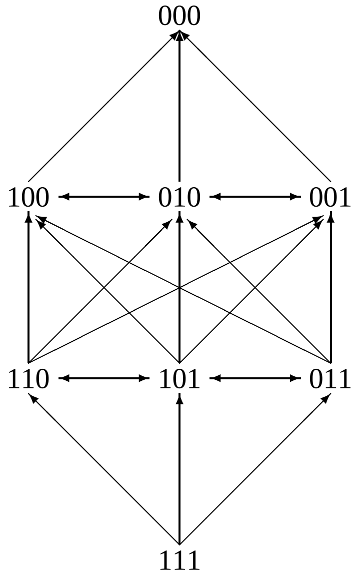

Figure 1. A graphical representation of the B-induced faithful total preorder $\preceq$, with $B = \{p, q, r\}$. For every $x, y \in U$, $x \preceq y$ iff (x, y) is in the reflexive transitive closure of the relation determined by the arrows.

reason the weakened versions of p and q should not contain any “part” of r that is not already contained in p or q. In fact, $p \vee q$ (or any wff logically equivalent to it) would be a suitable choice for the weakened versions of both p and q. Unfortunately, the proposal mentioned above does not give the desired weakened versions of p and q (although it does ensure that the resulting infobase contains r). Since $\text{Min}_{\preceq}(\neg(p \wedge q)) = \{101, 011\}$, the weakened version of p would be logically equivalent to $p \vee (q \wedge r)$ and the weakened version of q would be logically equivalent to $q \vee (p \wedge r)$.

The problem in the example above seems to be that r has too great an influence on the weakened versions of p and q. When deciding which interpretations to add to the models of p (or q) to obtain its weakened version, we should ensure that any undue influence of r is neutralised. Accordingly, we should not just add the minimal models of $\neg(p \wedge q)$, but also any other models of $\neg(p \wedge q)$ that behave exactly like the minimal models with respect to the wffs p and q, but that might differ from the minimal models on the truth value of r.

DEFINITION 5. For $Z \subseteq L$ and $x, y \in U$, x is *Z-equivalent* to y, written $x \equiv_Z y$, iff for every $\alpha \in Z$, $x \Vdash \alpha$ iff $y \Vdash \alpha$.

In the example above, 100 and 010 are $\{p, q\}$-equivalent to the minimal models 101 and 011 respectively, and adding them to the models of p (and q) as well results in weakened versions of p and q that are indeed logically equivalent to $p \vee q$.

In general, we obtain the weakened version of every ϕ-discarded wff ψ as follows. We need some appropriate set of interpretations that can be added to the models of ψ to obtain the set of models of its weakened version. We use the set of minimal models of $\neg\phi$ as our starting point and then try to expand it so that only the ϕ-discarded wffs have any influence, thus neutralising the possible influence of any of the ϕ-retained wffs. This is accomplished by including all the models of $\neg\phi$ that are $B^{-\phi}$-equivalent to some minimal model of $\neg\phi$.

DEFINITION 6. Let $\preceq$ be the B-induced faithful total preorder. For every $x \in \mathrm{Min}_{\preceq}(\neg\phi)$, we let $N_x(\neg\phi) = \{y \in M(\neg\phi) \mid y \equiv_{B^{-\phi}} x\}$, and we let $N_B(\neg\phi) = \bigcup_{x\in\mathrm{Min}_{\preceq}(\neg\phi)} N_x(\neg\phi)$. We refer to $N_B(\neg\phi)$ as the *ϕ-neutralised models* of B.

We take the ϕ-neutralised models as the set of interpretations to be added to the models of each ϕ-discarded wff. We can think of the ϕ-neutralised models as a set of interpretations in which the influence of the ϕ-retained wffs has been removed, but in which the ϕ-discarded wffs have the same impact as on the minimal models of $\neg\phi$. Of course, such a definition only makes sense if these sets of interpretations can be axiomatised by single wffs. While this is immediate for the finitely generated propositional logics, the next result shows that it also holds in the more general case.

DEFINITION 7. Let $\preceq$ be the B-induced faithful total preorder, let $-$ be the B-induced contraction function, and let $B^{-\phi}_{\preceq} = \{C \subseteq B \mid \exists x \in \mathrm{Min}_{\preceq}(\neg\phi) \cap M(C)$ s.t. $|C| = x_B\}$ (where x_B is the B-number of x). For every ϕ-discarded wff ψ, we define the *ϕ-weakened version* of ψ as

$$w_B^{-\phi}(\psi) = \begin{cases} \psi \vee \left(\bigvee_{C\in B^{-\phi}_{\preceq}} \left(\bigwedge\left[C \setminus \left(B \setminus B^{-\phi}\right)\right]\right) \wedge \left(\bigwedge \neg(B^{-\phi} \setminus C)\right) \wedge \neg\phi\right) \\ \qquad \text{if} B^{-\phi}_{\preceq} \neq \emptyset, \\ \psi \vee \neg\phi \\ \qquad \text{otherwise.} \end{cases}$$

PROPOSITION 1. *Let $\preceq$ be the B-induced faithful total preorder. For every ϕ and every ϕ-discarded wff ψ, $M(w_B^{-\phi}(\psi)) = M(\psi) \cup N_B(\neg\phi)$.*

Proof. Define $B^{-\phi}_{\preceq}$ as in Definition 7. We only consider the case where $B^{-\phi}_{\preceq} \neq \emptyset$. Then every $x \in \mathrm{Min}_{\preceq}(\neg\phi)$ is a model of some $C \in B^{-\phi}_{\preceq}$. Pick any $C \in B^{-\phi}_{\preceq}$, and any $x \in \mathrm{Min}_{\preceq}(\neg\phi) \cap M(C)$. Observe that every model of $C \cup \{\neg\phi\}$ is a minimal element of $M(\neg\phi)$, which ensures that every element of $B^{-\phi} \setminus C$ is false in all the models of $C \cup \{\neg\phi\}$. We record this result formally.

$$\forall\gamma \in B^{-\phi} \setminus C, \quad \forall y \in M(C \cup \{\neg\phi\}), \quad y \not\Vdash \gamma. \tag{1}$$

We show that $M\left(\left[C \setminus (B \setminus B^{-\phi})\right] \cup \neg(B^{-\phi} \setminus C) \cup \{\neg\phi\}\right) = N_x(\neg\phi)$. From Equation (1) it follows that $x \not\Vdash \gamma$ for every $\gamma \in B^{-\phi} \setminus C$ and therefore that $x \in M\left(\left[C \setminus (B \setminus B^{-\phi})\right] \cup \neg(B^{-\phi} \setminus C) \cup \{\neg\phi\}\right)$. Now pick any $y \in$

$M\left(\left[C \setminus (B \setminus B^{-\phi})\right] \cup \neg(B^{-\phi} \setminus C) \cup \{\neg\phi\}\right)$ and any $\beta \in B^{-\phi}$. If $\beta \in C$ then clearly $x \Vdash \beta$ iff $y \Vdash \beta$, so suppose $\beta \notin C$. Then by Equation (1) again, $x \nVdash \beta$. Furthermore, since $y \in M(\neg(B^{-\phi} \setminus C))$, it follows that $y \nVdash \beta$ and thus that $x \Vdash \beta$ iff $y \Vdash \beta$. Finally, it is clear that $y \in M(\neg\phi)$, and thus $y \in N_x(\neg\phi)$. Conversely, pick any $y \in N_x(\neg\phi)$. Clearly, $y \in M(\neg\phi)$, and since $x \in M\left(\left[C \setminus (B \setminus B^{-\phi})\right] \cup \neg(B^{-\phi} \setminus C) \cup \{\neg\phi\}\right)$, so is y.

It is clear that $M\left(\left[C \setminus (B \setminus B^{-\phi})\right] \cup \neg(B^{-\phi} \setminus C) \cup \{\neg\phi\}\right)$ is axiomatised by the wff

$$(\neg\phi)^C = \left(\bigwedge \left[C \setminus (B \setminus B^{-\phi})\right]\right) \wedge \left(\bigwedge \neg(B^{-\phi} \setminus C)\right) \wedge \neg\phi$$

and it thus follows that $M((\neg\phi)^C) = N_x(\neg\phi)$. So we have shown that if $B_{\preceq}^{-\phi} \neq \emptyset$, then

$$\forall\, C \in B_{\preceq}^{-\phi}, \exists x \in \mathrm{Min}_{\preceq}(\neg\phi) \cap M(C) \quad \text{and} \tag{2}$$

$$\forall\, C \in B_{\preceq}^{-\phi}, \forall x \in \mathrm{Min}_{\preceq}(\neg\phi) \cap M(C), M\left((\neg\phi)^C\right) = N_x(\neg\phi). \tag{3}$$

We now show that $N_B(\neg\phi) = M\left(\bigvee_{C \in B_{\preceq}^{-\phi}} (\neg\phi)^C\right)$, from which the required result follows. Pick a $y \in N_B(\neg\phi)$. There is an $x \in \mathrm{Min}_{\preceq}(\neg\phi)$ such that $y \in N_x(\neg\phi)$, and by Equation (3) it follows that for some $C \in B_{\preceq}^{-\phi}$, $y \in N_x(\neg\phi) = M\left((\neg\phi)^C\right)$. So clearly $y \in M\left(\bigvee_{C \in B_{\preceq}^{-\phi}} (\neg\phi)^C\right)$. Conversely, pick any $y \in M\left(\bigvee_{C \in B_{\preceq}^{-\phi}} (\neg\phi)^C\right)$. Then y is a model of $(\neg\phi)^C$ for some $C \in B_{\preceq}^{-\phi}$. By Equation (2) there is an $x \in \mathrm{Min}_{\preceq}(\neg\phi) \cap M(C)$, and by Equation (3), $N_x(\neg\phi) = M\left((\neg\phi)^C\right)$. So $y \in N_x(\neg\phi)$ and thus $y \in N_B(\neg\phi)$. □

We are now in a position to give a formal definition of infobase contraction.

DEFINITION 8. Let $\preceq$ be the B-induced faithful total preorder, and let $-$ be the B-induced contraction function. The *infobase B-contraction function* is defined as

$$B \ominus \phi = (B \setminus B^{-\phi}) \cup \{w_B^{-\phi}(\psi) \mid \psi \in B^{-\phi}\}.$$

We conclude this section with an example to illustrate the partial construction of an infobase B-contraction function for a propositional language generated by two atoms.

Example 2. Let $B = \{p, q\}$. Figure 2 contains a graphical representation of the B-induced faithful total preorder $\preceq$. We show how to construct $B \ominus p$ and $B \ominus (p \wedge q)$. Let $-$ be the B-induced contraction function. Then $Cn(B) - p = Cn(q)$, $B^{-p} = \{p\}$ and $B_{\preceq}^{-p} = \{\{q\}\}$. So $B \ominus p = \{w_B^{-p}(p), q\}$, where $w_B^{-p}(p) = p \vee (\top \wedge \neg p \wedge \neg p) \equiv \top$. Similarly,

$$Cn(B) - (p \wedge q) = Cn(p \vee q), \quad B^{-(p\wedge q)} = \{p, q\} \quad \text{and} \quad B_{\preceq}^{-p\wedge q} = \{\{p\}, \{q\}\}.$$

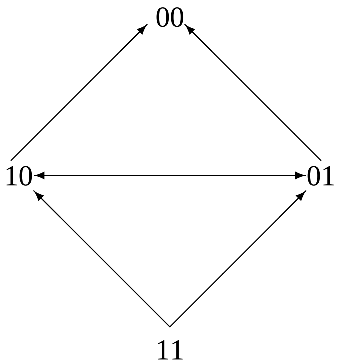

Figure 2. A graphical representation of the B-induced faithful total preorder $\preceq$, with $B = \{p, q\}$. For every $x, y \in U$, $x \preceq y$ iff (x, y) is in the reflexive transitive closure of the relation determined by the arrows.

So $B \ominus (p \wedge q) = \{w_B^{-p\wedge q}(p), w_B^{-p\wedge q}(q)\}$, where

$$w_B^{-p\wedge q}(p) = p \vee \left(\left[p \wedge \neg q \wedge \neg(p \wedge q)\right] \vee \left[q \wedge \neg p \wedge \neg(p \wedge q)\right]\right), \quad \text{and}$$
$$w_B^{-p\wedge q}(q) = q \vee \left(\left[p \wedge \neg q \wedge \neg(p \wedge q)\right] \vee \left[q \wedge \neg p \wedge \neg(p \wedge q)\right]\right).$$

It is readily verified that both $w_B^{-p\wedge q}(p)$ and $w_B^{-p\wedge q}(q)$ are logically equivalent to $p \vee q$.

3.1.1. *Required Properties of Infobase Contraction*

In the discussion of infobase contraction thus far it has been implied that

- the ϕ-weakened versions of the ϕ-discarded wffs are appropriate choices for weakened versions of these wffs, that
- the B-induced contraction function is the theory $Cn(B)$-contraction function associated with infobase B-contraction, and that
- infobase contraction is not concerned with the syntactic form of the wffs in an infobase.

The first point has already been dealt with in the previous section. This section is devoted to a justification of the remaining two claims. The first result presented here is a preliminary one, indicating that the models of the ϕ-retained wffs that are also ϕ-neutralised models, are precisely the minimal models of $\neg\phi$.

LEMMA 1. *If $\preceq$ is the B-induced faithful total preorder, then*

$$N_B(\neg\phi) \cap M(B \setminus B^{-\phi}) = \text{Min}_{\preceq}(\neg\phi).$$

Proof. Let $-$ be the B-induced contraction function. By definition, $B \setminus B^{-\phi} \subseteq Cn(B)-\phi$ and thus $M(B) \cup \mathrm{Min}_{\preceq}(\neg\phi) \subseteq M(B \setminus B^{-\phi})$. Furthermore, $\mathrm{Min}_{\preceq}(\neg\phi) \subseteq N_B(\neg\phi)$, and so $\mathrm{Min}_{\preceq}(\neg\phi) \subseteq N_B(\neg\phi) \cap M(B \setminus B^{-\phi})$. Conversely, pick any $y \in N_B(\neg\phi) \cap M(B \setminus B^{-\phi})$. That is, y satisfies all the ϕ-retained wffs, y is a model of $\neg\phi$ and there is a minimal model x of $\neg\phi$ that satisfies exactly the same ϕ-discarded wffs as y does. Because $x \in \mathrm{Min}_{\preceq}(\neg\phi)$ it follows from the definition of $-$ and $B^{-\phi}$ that x also satisfies all the wffs in $B \setminus B^{-\phi}$. So x and y satisfy exactly the same wffs in B, which means that $y \in \mathrm{Min}_{\preceq}(\neg\phi)$. □

The result above is used to prove that the B-induced contraction function $-$ is the $Cn(B)$-contraction function associated with the infobase B-contraction function.

PROPOSITION 2. *Let $\ominus$ be the infobase B-contraction function, and let $-$ be the B-induced contraction function. Then $Cn(B) - \phi = Cn(B \ominus \phi)$.*

Proof. Let $\preceq$ be the B-induced faithful total preorder. By Proposition 1,

$$\begin{aligned} M(B \ominus \phi) &= \left[\bigcap_{\psi \in B^{-\phi}} M\left(w_B^{-\phi}(\psi)\right)\right] \cap M(B \setminus B^{-\phi}) \\ &= \left[\bigcap_{\psi \in B^{-\phi}} (M(\psi) \cup N_B(\neg\phi))\right] \cap M(B \setminus B^{-\phi}) \\ &= \left[\left(\bigcap_{\psi \in B^{-\phi}} M(\psi)\right) \cup N_B(\neg\phi)\right] \cap M(B \setminus B^{-\phi}) \\ &= \left(M(B^{-\phi}) \cup N_B(\neg\phi)\right) \cap M(B \setminus B^{-\phi}) \\ &= M(B) \cup \left(N_B(\neg\phi) \cap M(B \setminus B^{-\phi})\right) \\ &= M(B) \cup \mathrm{Min}_{\preceq}(\neg\phi) \quad \text{by Lemma 1,} \end{aligned}$$

and thus $Cn(B) - \phi = Cn(B \ominus \phi)$. □

And next, we show that the syntactic form of wffs in an infobase is irrelevant.

DEFINITION 9. Two infobases B and C are *element-equivalent*, written as $B \approx C$, iff for every $\phi \in B$ such that $\nvDash \phi$, there is a logically equivalent wff $\psi \in C$, and for every $\psi \in C$ such that $\nvDash \psi$, there is a logically equivalent wff $\phi \in B$.

PROPOSITION 3. *If $B \approx C$ then $B \ominus \phi \approx C \ominus \phi$.*

Proof. Since B and C are element-equivalent, $x_B = x_C$ for every $x \in U$, and so the B-induced faithful total preorder is exactly the same as the C-induced faithful preorder. Therefore $B^{-\phi} \approx C^{-\phi}$, $B \setminus B^{-\phi} \approx C \setminus C^{-\phi}$, and thus $N_B(\neg\phi) =$

$N_C(\neg\phi)$. So, for every $\psi \in B^{-\phi}$ and every $\chi \in C^{-\phi}$ such that $\psi \equiv \chi$, $w_B^{-\phi}(\psi) \equiv w_C^{-\phi}(\chi)$ by Proposition 1, and for every $\psi \in C^{-\phi}$ and every $\chi \in B^{-\phi}$ such that $\psi \equiv \chi$, $w_C^{-\phi}(\psi) \equiv w_B^{-\phi}(\chi)$ by Proposition 1, from which the required result follows. □

3.1.2. *Infobase Contraction and Reason Maintenance*

Reason maintenance is a process in which the removal of a basic belief forces the removal of the consequences of the basic belief as well, unless they can be derived from other basic beliefs. In the context of infobase change, the wffs in an infobase B can be viewed as the basic beliefs of the belief set generated by B. Reason maintenance would thus ensure that the contraction of B by a wff ϕ in B results in the removal of all the wffs that are dependent on ϕ for being in $Cn(B)$. With the notion of B-dependence, Fuhrmann (1991) has given a precise meaning to the idea of a wff being dependent on ϕ for being in $Cn(B)$.

DEFINITION 10. A wff $\psi \in L$ is *B-dependent on ϕ* iff $\phi \in B$ and $\psi \in Cn(B)$, but $\psi \notin Cn(B \setminus \{\phi\})$.

The next result shows that infobase contraction incorporates reason maintenance.

PROPOSITION 4. *If ψ is B-dependent on ϕ then $\psi \notin Cn(B \ominus \phi)$.*

Proof. Since $\psi \in Cn(B)$, but $\psi \notin Cn(B \setminus \{\phi\})$, there has to be a model x of $B \setminus \{\phi\}$ in which both ϕ and ψ are false. So $x \in M(\neg\phi)$ and $x \notin M(B)$. Now, there is only one wff in B, namely ϕ, that is false in x, so any interpretation y for which $y_B > x_B$ has to be a model of B. Therefore $x \in \text{Min}_{\preceq}(\neg\phi)$, and because $x \notin M(\psi)$, it follows that $\psi \notin Cn(B) - \phi$, where $-$ is the B-induced contraction function. So $\psi \notin Cn(B \ominus \phi)$ by Proposition 2. □

Of course, the contraction of B by a wff ϕ in B is not the only way to remove ϕ from the infobase. In the light of this, it seems reasonable to inquire whether the wffs that are B-dependent on ϕ will also be discarded if ϕ is discarded during the contraction of B by some wff other than ϕ itself. That is, if $\phi \in B$ and $\phi \notin Cn(B \ominus \chi)$, will it be the case that $\psi \notin Cn(B \ominus \chi)$ for every ψ that is B-dependent on ϕ? This property is known as the *filtering condition* (Fuhrmann, 1991). It is easy to see that infobase contraction can violate the filtering condition. For example, it is readily verified that a contraction of the infobase $B = \{p \wedge q\}$ by p results in an infobase in which $p \wedge q$ is replaced by the wff $w_B^{-p}(p \wedge q)$ which is logically equivalent to $p \rightarrow q$. And since $w_B^{-p}(p \wedge q)$ is clearly B-dependent on $p \wedge q$, the filtering condition is violated. But such a violation is to be expected. Given the intuition associated with infobases, the filtering condition is clearly too strong a requirement to impose on infobase contraction. For the filtering condition requires that $Cn(B \ominus \chi) = Cn(\top)$ for any singleton infobase B, and any $\chi \in Cn(B)$

(where $\nvDash \chi$), thus leaving no room for weakening the wff in B to anything but a logically valid wff.

3.2. INFOBASE REVISION

Infobase revision is defined by appealing to the infobase analogue of the Levi Identity.

DEFINITION 11. The *infobase B-revision function* is defined as $B \circledast \phi = (B \ominus \neg\phi) \cup \{\phi\}$, where $\ominus$ is the infobase B-contraction function.

Given the connection between infobase contraction and revision, it is to be expected that infobase revision satisfies properties that are similar to those proved in Sections 3.1.1 and 3.1.2. The next corollary shows that this is indeed the case.

DEFINITION 12. *The B*-induced revision function $*$ *is defined as*

$$Cn(B) * \phi = Th(\mathrm{Min}_{\preceq}(\phi)),$$

where $\preceq$ is the B-induced faithful total preorder.

COROLLARY 1. *Let $\circledast$ be the infobase B-revision function and let $*$ be the B-induced revision function.*

1. *If $B \approx C$ then $B \circledast \phi \approx C \circledast \phi$.*
2. *$Cn(B \circledast \phi) = Cn(B) * \phi$.*
3. *If ψ is B-dependent on ϕ, then $\psi \notin Cn(B \circledast \neg\phi)$.*

Proof.

1. Follows from Proposition 3.
2. Follows from Proposition 2, and by noting that $\mathrm{Min}_{\preceq}(\phi) \subseteq M(B)$ if $\neg\phi \notin Cn(B)$, and that $Cn(B)*\phi = Th(\mathrm{Min}_{\preceq}(\phi))$, where $\preceq$ is the B-induced faithful total preorder.
3. Follows from part (2) of this corollary, and by an argument similar to the proof of Proposition 4. □

Corollary 1 thus shows that infobase revision is insensitive to the syntactic form of the wffs in an infobase, that the theory revision function associated with the infobase revision function is the B-induced revision function, and that infobase revision can be said to perform reason maintenance. It is also possible to describe infobase revision directly, and not in terms of infobase contraction. Let $B^{*\phi} = \{\psi \in B \mid \psi \notin Cn(B) * \phi\}$, where $*$ be the B-induced revision function. Then it is easy to see that $B^{-\neg\phi}$ and $B^{*\phi}$ are identical.

LEMMA 2. $B^{-\neg\phi} = B^{*\phi}$.

Proof. $\psi \in B^{-\neg\phi}$ iff $M(B) \subseteq M(\psi)$ and $\text{Min}_{\preceq}(\phi) \not\subseteq M(\psi)$, iff $\psi \notin B^{*\phi}$. □

Setting $w_B^{*\phi}(\psi)$ equal to $w_B^{-\neg\phi}(\psi)$, this result enables us to define infobase revision directly.

PROPOSITION 5. *Let $\preceq$ be the B-induced faithful total preorder and let $*$ be the B-induced revision function. Then the* infobase B-revision function $\circledast$ *can also be described as*

$$B \circledast \phi = B^{*\phi} \cup \{w_B^{*\phi}(\psi) \mid \psi \in B \setminus B^{*\phi}\} \cup \{\phi\}.$$

Proof. Follows from Lemma 2. □

To conclude this section, we provide an infobase change analogue of the Harper Identity.

PROPOSITION 6. *The infobase B-contraction function $\ominus$ can be described in terms of the infobase B-revision function $\circledast$ as*

$$B \ominus \phi \approx B \ominus \neg\neg\phi = \begin{cases} (B \circledast \neg\phi) \setminus \{\neg\phi\} & \textit{if } \neg\phi \notin B, \\ B \circledast \neg\phi, & \textit{otherwise}. \end{cases}$$

Proof. If $\phi \notin Cn(B)$ then $B \ominus \neg\neg\phi = B$, and the result easily follows. If $\phi \in Cn(B)$, the result follows by noting that if $\neg\phi \in B$ then $\neg\phi \in B \ominus \neg\neg\phi$. □

4. Related Approaches to Base Change

Infobase change relies heavily on the faithful total preorder obtained by counting the number of non-equivalent wffs in a base B. As such, its roots can be found in the work of Dalal (1988), Borgida (1985), Satoh (1988), Weber (1986) and Winslett (1988), all of which use the idea of distinguishing between interpretations based on the number of propositional atoms that they satisfy (at least in the propositional case). However, these approaches do not distinguish between different bases generating the same belief set, and are thus more properly classified as instances of theory change.

Unlike infobase contraction, most versions of base contraction require the base resulting from a base contraction operation to be a subset of the original base. Two notable exceptions are the base contraction operations of Nebel (1989, 1990, 1991, 1992), and Nayak (1994), which allow wffs into the resulting base that were not in the original base. In this section we compare these two approaches with infobase change.

4.1. NEBEL'S APPROACH

Nebel's base change operations (1990, 1991, 1992) make use of an *epistemic relevance* ordering on the wffs in the belief set generated by the base, which is taken to denote relative epistemic importance. This is a generalisation of the case considered in Nebel (1989), which can be seen as the special case where all wffs in the base have equal epistemic weight. Since the latter is closely related to infobase change, we shall mainly concern ourselves with the work in Nebel (1989).

Nebel's construction of base contraction functions uses the maximal subsets of a set X that do not entail ϕ. It can thus be seen as a generalisation of the construction of the partial meet functions of Alchourrón et al. (1985). For every $X \subseteq L$, let $X \downarrow \phi$, the set of *remainders of* X *after removing* ϕ, be defined as

$$X \downarrow \phi = \{Y \subseteq X \mid Y \not\models \phi \text{ and for every } Z \subseteq L \text{ such that } Y \subseteq Z \subset X, Z \models \phi\}.$$

Nebel defines the base contraction function $\hat{\sim}$, in a somewhat opaque fashion, as

$$B \hat{\sim} \phi = \begin{cases} \bigvee_{C \in (B \downarrow \phi)} C \wedge (B \vee \{\neg\phi\}) & \text{if } \not\models \phi, \\ B & \text{otherwise.} \end{cases}$$

This construction is justified by a closer look at the theory contraction function associated with $\hat{\sim}$. He defines a B-faithful weak partial order $\leq$ as: $x \leq y$ iff $(Th(x) \cap B) \supseteq (Th(y) \cap B)$, and then obtains a $Cn(B)$-contraction function $\hat{-}$ from $\leq$ as follows:*

$$Cn(B) \hat{-} \phi = Th(M(B) \cup \text{Min}_{\leq}(\neg\phi)).$$

He then proceeds to show that $\hat{-}$ is the $Cn(B)$-contraction function associated with $\hat{\sim}$ (i.e., $Cn(B) \hat{-} \phi = Cn(B \hat{\sim} \phi)$), and that $\hat{-}$ satisfies $(K-1)$ to $(K-7)$, but does not, in general, satisfy $(K-8)$.

A comparison of Nebel's $Cn(B)$-contraction function $\hat{-}$ (which is obtained from $\leq$) with the B-induced contraction function shows that the intuitions employed in both cases are very similar. But whereas $\leq$ is defined in terms of the satisfaction of maximal *subsets* of B, the corresponding B-induced faithful total preorder relies on the satisfaction of the maximum *number* of wffs in B. While this difference allows for Nebel's $\hat{-}$ to be defined for infinite bases as well, it ensures that $\hat{-}$ does not always satisfy $(K-8)$, while the B-induced contraction function does. Below we provide an example in which it seems desirable for a base contraction operation to satisfy $(K-8)$, at least under the assumption of the independence of the wffs in a base B.

Example 3. Let $B = \{p \vee q, \neg p \vee q, p\}$ and let $\sim$ be a base contraction function in which the wffs in B are regarded as being independently obtained. A contraction with $p \wedge q$ would force us to remove at least one of p and q, and since

* Nebel's construction of the theory contraction function $\hat{-}$ is phrased in terms of partial meet functions, but it is easily seen that it can also be phrased semantically, as we have done.

$p \in B$ but $q \notin B$, it seems reasonable to require that if one of the two is retained, it should be p and not q. So, regardless of whether p is being retained, q should not be an element of $Cn(B \sim (p \wedge q))$. Furthermore, since $p \vee q$ is explicitly contained in B, a contraction of B by $p \wedge q$ should not remove $p \vee q$, and we should thus have $p \vee q \in Cn(B \sim (p \wedge q))$. Finally, although the presence of both $p \vee q$ and $\neg p \vee q$ in B suggests that p and q are independent (since $p \vee q$ is logically equivalent to $\neg p \rightarrow q$, and $\neg p \vee q$ to $p \rightarrow q$), this is, to some extent, offset by the presence in B of both p and $\neg p \vee q$. The inconclusive evidence regarding the independence of p and q, coupled with the fact that p itself is in B, then suggests that p should be an element of $Cn(B \sim q)$. It is easy to see that the failure of the intuition expressed above would amount to a violation of $(K-8)$. By taking ϕ as p and ψ as q, it is easily seen that Nebel's $Cn(B)$-contraction function $\hat{-}$ violates $(K-8)$ ($p \vee q \in Cn(B)\hat{-}(p \wedge q)$, but $p \vee q \notin Cn(B)\hat{-}q$).

Nebel also considers a modification of $\hat{-}$ that satisfies $(K-8)$ (which allows him to set $B\hat{\sim}\phi$ equal to some element of $B \downarrow \phi$) but it presupposes a linear order on the wffs in B, which is a very strong restriction indeed. The restriction is relaxed to a total preorder in Nebel (1990, 1991, 1992), but then $(K-8)$ does not hold in the general case.

We have thus far considered the $Cn(B)$-contraction function $\hat{-}$ in detail, but have said very little about $\hat{\sim}$ itself. From some comments made in his conclusion, it seems that Nebel regards the set of wffs $B\hat{\sim}\phi$ merely as a convenient finite representation from which the belief set $B\hat{-}\phi$ can be generated, and nothing more. He writes: "…iterated contractions were ignored because they present serious problems," and "Choosing the 'right' form of the premises seems to be one of the central tasks before any kind of belief revision can be applied." The latter statement seems to suggest that $B\hat{\sim}\phi$ cannot be seen as a base with the wffs contained in it being epistemologically more important than the wffs in $Cn(B\hat{\sim}\phi)$, a view that is also supported by his proposal for a base revision function $\hat{*}$. He defines $B\hat{*}\phi$ as $(B\hat{\sim}\neg\phi) \wedge \{\phi\}$, which means that the newly obtained basic belief ϕ occurs in $B\hat{*}\phi$ as a conjunct of every wff in $(B\hat{\sim}\neg\phi)$. And there certainly is no intuition of a weakening of the wffs contained in B, as with infobase change. For example, if $B = \{p, q, r\}$, it can be verified that $B\hat{\sim}(p \wedge q \wedge r)$ contains 24 elements and is element-equivalent to the set $\{p \vee q, p \vee r, q \vee r, p \vee q \vee r\}$. In contrast, $B \ominus (p \wedge q \wedge r)$ (where $\ominus$ is the infobase B-contraction function) contains three logically non-equivalent wffs (weakened versions of each of the wffs in B) and is element-equivalent to the set $\{p \vee (q \wedge r), q \vee (p \wedge r), r \vee (p \wedge q)\}$.

4.2. NAYAK'S APPROACH

Nayak's (1994) approach to base change is more general than infobase contraction since it accommodates infinite bases and allows multiple possible outcomes. It takes Fuhrmann's (1991) generalised safe contraction as a starting point. When contracting B by ϕ (a base contraction function that we denote by $\ddot{\sim}$) he first finds the set $E(\phi)$ of minimal subsets of B that entail ϕ. The idea is to construct a *reject set* $R(\phi)$ (wffs of B that will be discarded), consisting of wffs from every element of $E(\phi)$. To ensure that the $Cn(B)$-contraction function associated with $\ddot{\sim}$ satisfies the first six AGM contraction postulates, except for $(K-5)$, he assumes a *choice function* C from $\wp B$ to $\wp B$ that picks the "most rejectable" elements of any subset of B. Up to this point the construction corresponds roughly to Fuhrmann's base contraction. However, Fuhrmann's version of the choice function does not have to conform to the stringent restrictions that Nayak places on C. Furthermore, Nayak does not take the set $R_0(\phi)$, which consists of the most rejectable elements of all members of $E(\phi)$, to be the reject set, as Fuhrmann does. Instead, he uses C to choose a particular subset of $R_0(\phi)$, which also happens to be an element of $B \downarrow \phi$, as the reject set $R(\phi)$. $B\ddot{\sim}\phi$ is then defined as the wffs in B that are not rejected, together with weakened versions of the rejected wffs. To be precise, $B\ddot{\sim}\phi = B \setminus R(\phi) \cup \{\psi \to \phi \mid \psi \in R(\phi)\}$. Nayak proves that the $Cn(B)$-contraction function $\ddot{-}$ associated with $\ddot{\sim}$ satisfies all eight AGM contraction postulates. The addition of the weakened versions of wffs in the reject set ensures that $\ddot{-}$ satisfies $(K-5)$, but it is currently unclear whether it plays a role in the satisfaction of $(K-7)$ and $(K-8)$ as well.

The strict conditions imposed on C, together with the insistence that the reject set $R(\phi)$ be an element of $B \downarrow \phi$, are akin to placing a linear order on B. This means that Nayak's base contraction function $\ddot{\sim}$ is closely related to Nebel's modified version of the base contraction function $\hat{\sim}$, for which $B\hat{\sim}\phi$ is an element of $B \downarrow \phi$. It is thus difficult to draw a direct comparison between $\ddot{\sim}$ and infobase contraction, mainly because the construction of $\ddot{\sim}$ needs so much more extra-logical information. A feature that these two forms of base contraction do have in common concerns the wffs contained in the resulting base after a contraction has taken place. Both retain a number of wffs in B and replace the wffs that are removed with weakened versions. Currently, the closest we can come to a comparison is to give an example showing that versions of $\ddot{\sim}$ that cater for situations in which less extra-logical information is available will probably not always give the desired results, at least not when the wffs in a base are assumed to be independent.

Example 4. Let $B = \{p, q\}$. The requirement that the reject set be a subset of B seems to form an integral part of Nayak's approach, which means that the reject set $R(p \wedge q)$ has to be one of $\emptyset$, $\{p\}$ or $\{q\}$, irrespective of any restrictions on the choice function C. The only candidates for $B\ddot{\sim}(p \wedge q)$ are thus $\{(p \wedge q) \to p, (p \wedge q) \to q\}$, $\{p, (p \wedge q) \to q\}$ and $\{q, (p \wedge q) \to p\}$. Now, if p and q have equal weight then the desired result when contracting B with $p \wedge q$ is $\{p \vee q\}$ or some element-equivalent set, a set of wffs which Nayak's approach is not able to produce. This is in contrast with infobase contraction, which yields $B \ominus (p \wedge q) \approx \{p \vee q\}$, as we have seen in Example 2.

5. Iterated Infobase Change

Since the revision and contraction of an infobase B both yield unique infobases, it is possible to perform *iterated* infobase change. It is thus worth investigating to what extent iterated infobase revision fits into frameworks proposed for iterated revision. In this section we focus on the proposals of Lehmann (1995) and Darwiche and Pearl (1997).

Darwiche and Pearl extend the AGM framework to be able to accommodate iterated revision for finitely generated propositional logics. To do so, they find it necessary to perform revision on the level of epistemic states, where an epistemic state is an ordered pair $\Phi = (K_\Phi, \preceq_\Phi)$ such that $\preceq_\Phi$ is a K_Φ-faithful total preorder. The AGM revision postulates are modified accordingly and, with the aim of constraining the possible ways of performing iterated revision, Darwiche and Pearl then propose four additional postulates (the DP-postulates).

If we are to measure iterated infobase revision against the DP-postulates, we need to go a step further and describe revision at the level of infobases. This is necessary since infobases contain more information than epistemic states. That is, while every infobase B is uniquely associated with the epistemic state $(Cn(B), \preceq)$, where $\preceq$ is the B-induced faithful total preorder, the same epistemic state may be associated with different infobases. For example, letting $B = \{p, q\}$ and $C = \{p \wedge q, p \vee q\}$, it is easy to check that the B-induced faithful total preorder and the C-induced faithful total preorder are identical. Furthermore, the fact that we only deal with finitely generated propositional logics makes it easy to see that every epistemic state is obtained from some infobase.

LEMMA 3. *For every epistemic state $\Phi = (K_\Phi, \preceq_\Phi)$ there is an infobase B such that $K_\Phi = Cn(B)$ and $\preceq_\Phi$ is the B-induced faithful preorder.*

Proof. Pick any epistemic state $\Phi = (K_\Phi, \preceq_\Phi)$. Since L is a finitely generated propositional language, U contains a finite number of interpretations. The total preorder $\preceq_\Phi$ thus partitions U into a finite number of subsets (blocks). Let us assume that there are n such blocks. We assign each of them a unique index from 1 to n according to their relative positions in $\preceq_\Phi$, leaving us with the n indexed blocks $P_1, \ldots, P_n$. That is, for $1 \leq i, j \leq n$, $i < j$ iff for every $x \in P_i$ and every $y \in P_j$, $x \prec_\Phi y$. Now, for any $V \subseteq U$, let $f(V)$ denote some wff that axiomatises V. (Since L is finitely generated, such a wff always exists.) For $1 \leq i \leq n$, let $\phi_i \equiv f(\bigcup_{1 \leq j \leq i} P_j)$. We define an infobase B as follows: if $\bot \in K_\Phi$, then $B = \{\bot\} \cup \bigcup_{1 \leq i \leq n} \{\phi_i\}$, otherwise $B = \bigcup_{1 \leq i \leq n} \{\phi_i\}$. It is easily verified that $Cn(B) = K_\Phi$ and that $\preceq_\Phi$ is the B-induced faithful total preorder. □

More importantly, perhaps, is the fact that the extra information contained in infobases plays an important role in the process of revision (and contraction), as the next example shows.

Example 5. Let $B = \{p, q\}$ and $C = \{p \wedge q, p, q, p \vee q, p \rightarrow q, q \rightarrow p\}$. Clearly, $Cn(B) = Cn(C)$ and it is also easy to see that the B-induced faithful

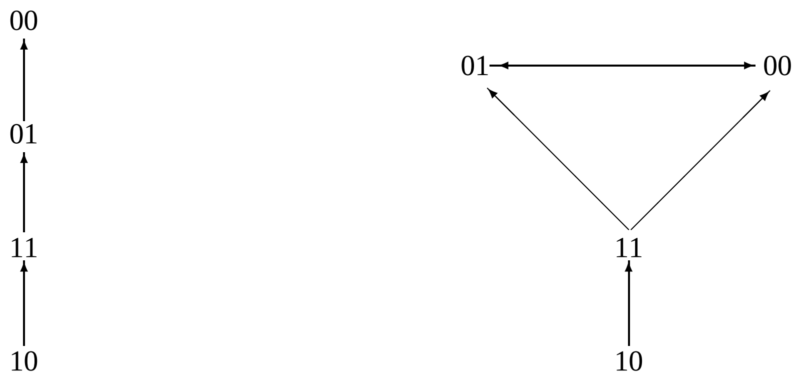

Figure 3. A graphical representation of the total preorders used in Example 5. On the left is the $[B \circledast (p \wedge \neg q)]$-induced faithful total preorder and on the right the $[C \circledast (p \wedge \neg q)]$-induced faithful total preorder. As usual, the applicable preorder is the reflexive transitive closure of the relation determined by the arrows.

total preorder and the C-induced faithful total preorder are identical, and are represented graphically in Figure 2. Furthermore, it can be verified that $B \circledast (p \wedge \neg q) \approx \{p, p \vee q, p \wedge \neg q\}$ and $C \circledast (p \wedge \neg q) \approx \{p, p \vee q, q \rightarrow p, p \wedge \neg q\}$. So $B \circledast (p \wedge \neg q)$ and $C \circledast (p \wedge \neg q)$ induce different faithful total preorders, as can be seen in Figure 3, even though the epistemic states obtained from B and C are identical.

To accommodate the richer structure of infobases, we provide versions of the DP-postulates on the level of infobases.

$(DP1)$ If $\psi \vDash \phi$ then $Cn((B \circledast \phi) \circledast \psi) = Cn(B \circledast \psi)$.

$(DP2)$ If $\psi \vDash \neg\phi$ then $Cn((B \circledast \phi) \circledast \psi) = Cn(B \circledast \psi)$.

$(DP3)$ If $\phi \in Cn(B \circledast \psi)$ then $\phi \in Cn((B \circledast \phi) \circledast \psi)$.

$(DP4)$ If $\neg\phi \notin Cn(B \circledast \psi)$ then $\neg\phi \notin Cn((B \circledast \phi) \circledast \psi)$.

It turns out that, in general, infobase revision satisfies none of these postulates, as the following example shows.

Example 6.

1. Let $B = \{p \leftrightarrow q, p \vee \neg q, \neg p \vee \neg q, \neg q\}$. It can be verified that

$$B \circledast (p \vee q) \approx \{p \vee \neg q, \neg p \vee \neg q, \neg q, p \vee q\},$$
$$Cn((B \circledast (p \vee q)) \circledast q) = Cn(q), \quad \text{and}$$
$$Cn(B \circledast q) = Cn(p \wedge q).$$

So $q \vDash p \vee q$, but $Cn((B \circledast (p \vee q)) \circledast q) \neq Cn(B \circledast q)$, which is a violation of $(DP1)$.

2. Let $B = \{p \vee q, q \rightarrow p, q\}$. It can be verified that

$$B \circledast p \wedge \neg q \approx \{p \vee q, q \rightarrow p, p \wedge \neg q\},$$
$$Cn((B \circledast p \wedge \neg q) \circledast \neg p) = Cn(\neg p), \quad \text{and}$$
$$Cn(B \circledast \neg p) = Cn(\neg p \wedge q).$$

So $\neg p \vDash \neg(p \wedge \neg q)$ but $Cn((B \circledast p \wedge \neg q) \circledast \neg p) \neq Cn(B \circledast \neg p)$; a violation of $(DP2)$.

3. Let $B = \{(\neg p \vee \neg q) \wedge r, (p \vee r) \wedge \neg q, (\neg p \vee r) \wedge \neg q, \neg(p \wedge \neg q \wedge r), \neg(p \wedge r), \neg(p \wedge (q \vee r))\}$. Now, let $\phi = ((q \leftrightarrow r) \rightarrow p) \rightarrow (p \wedge \neg q)$ and let $\psi = p \wedge \neg(q \leftrightarrow r)$. It then readily follows that

$$B \circledast \phi \approx \{\neg((p \vee \neg r) \wedge q), \neg(p \wedge \neg q \wedge r), \neg(p \wedge r), \neg(p \wedge (q \vee r)), \phi\},$$
$$Cn((B \circledast \phi) \circledast \psi) = Cn(\psi), \quad \text{and}$$
$$Cn(B \circledast \psi) = Cn(p \wedge \neg q \wedge r).$$

So $\phi \in Cn(B \circledast \psi)$, but $\phi \notin Cn((B \circledast \phi) \circledast \psi)$, which is a violation of $(DP3)$. Furthermore, let $\chi = (p \wedge \neg q \wedge r) \vee (\neg p \wedge q \wedge \neg r)$. It can be verified that

$$Cn((B \circledast \phi) \circledast \chi) = Cn(\neg p \wedge q \wedge \neg r), \quad \text{and}$$
$$Cn(B \circledast \chi) = Cn(\chi).$$

So $\neg\phi \notin Cn(B \circledast \chi)$, but $\neg\phi \in Cn((B \circledast \phi) \circledast \chi)$; a violation of $(DP4)$.

In lieu of these seemingly discouraging results, some explanatory comments are perhaps in order. A careful examination of the examples above shows that, with the exception of the violation of $(DP1)$, the counterexamples hinge on the logical equivalence of the weakened versions of some of the logically non-equivalent members of the infobase. For example, when revising the infobase $B = \{p, q\}$ with the wff $\neg(p \leftrightarrow q)$, the weakened versions of p and q are both logically equivalent to $p \vee q$. In turns out that this is the cause of the failure to satisfy the last three DP-postulates. In fact, it can be shown that a modified version of infobase revision in which the B-induced faithful total preorders take all wffs in the infobase into account, even if there are logically equivalent ones, satisfies the last three DP-postulates. We shall leave a description of this modified version of infobase change to future research.

Like Darwiche and Pearl, Lehmann assumes L to be finitely generated. He does not extend the AGM revision postulates, opting instead for a set of postulates describing arbitrary finite sequences of revisions. However, he does show that for a fixed belief set, his version of revision satisfies all the AGM revision postulates.

On the level of infobase revision, Lehmann's postulates can be reformulated as follows:*

$(L1)$ $Cn(B \circledast \phi)$ is satisfiable.

$(L2)$ $\phi \in Cn(B \circledast \phi)$.

$(L3)$ If $\psi \in Cn(B \circledast \phi)$ then $\phi \rightarrow \psi \in Cn(B)$.

$(L4)$ If $\phi \in Cn(B)$ then $Cn(B \circledast \psi) = Cn((B \circledast \phi) \circledast \psi)$.

$(L5)$ If $\psi \vDash \phi$ then $Cn(((B \circledast \phi) \circledast \psi) \circledast \chi) = Cn((B \circledast \psi) \circledast \chi)$.

$(L6)$ If $\neg\psi \notin Cn(B \circledast \phi)$ then $Cn(((B \circledast \phi) \circledast \psi) \circledast \chi) = Cn(((B \circledast \phi) \circledast \phi \wedge \psi) \circledast \chi)$.

$(L7)$ $Cn((B \circledast \neg\phi) \circledast \phi) \subseteq Cn(B) + \phi$.

Under Lehmann's assumption that we may only revise with satisfiable wffs, $(L1)$ corresponds to $(K*1)$ and $(K*6)$. Furthermore, $(L2)$ corresponds to $(K*2)$, and by observing that the assertion $\phi \rightarrow \psi \in Cn(B)$ is equivalent to $\psi \in Cn(B) + \phi$, we see that $(L3)$ corresponds to $(K*3)$. It is thus easily verified that infobase revision satisfies these three postulates. It does not satisfy the next three postulates, though, as the following example shows.

Example 7.

1. Let $B = \{p \wedge \neg q, p \vee q\}$. Clearly, $B \circledast p = \{p \wedge \neg q, p \vee q, p\}$. It can then be verified that $Cn((B \circledast p) \circledast q) = Cn(p \wedge q)$, but that $Cn(B \circledast q) = Cn(q)$. Taking p as ϕ and q as ψ, this is a violation of $(L4)$.
2. Let $B = \{p \leftrightarrow q, p \vee \neg q, \neg p \vee \neg q, \neg q\}$. It can be verified that

$$
\begin{aligned}
&B \circledast q \approx \{p \leftrightarrow q, p \vee \neg q, q\},\\
&B \circledast p \vee q \approx \{p \vee \neg q, \neg p \vee \neg q, \neg q, p \vee q\},\\
&(B \circledast p \vee q) \circledast q \approx \{p \vee q, q\},\\
&Cn(((B \circledast p \vee q) \circledast q) \circledast \neg q) = Cn(p \wedge \neg q), \quad \text{and}\\
&Cn((B \circledast q) \circledast \neg q) = Cn(\neg p \wedge \neg q).
\end{aligned}
$$

   Taking $p \vee q$ as ϕ, q as ψ, and $\neg q$ as χ, this constitutes a violation of $(L5)$.
3. Let $B = \{p \vee q, p \vee \neg q\}$. Clearly,

$$
\begin{aligned}
&B \circledast p = \{p \vee q, p \vee \neg q, p\},\\
&(B \circledast p) \circledast q = \{p \vee q, p \vee \neg q, p, q\}, \quad \text{and}\\
&(B \circledast p) \circledast p \wedge q = \{p \vee q, p \vee \neg q, p, p \wedge q\}.
\end{aligned}
$$

* Actually, $(L4)$ to $(L6)$ are weakened versions of Lehmann's corresponding three postulates, the latter dealing with sequences of revisions. We show below that infobase revision does not satisfy these three postulates. Consequently, it will not satisfy Lehmann's original postulates either.

It can be verified that

$$Cn(((B \circledast p) \circledast q) \circledast \neg p) = Cn(\neg p \wedge q), \quad \text{and}$$
$$Cn(((B \circledast p) \circledast p \wedge q) \circledast \neg p) = Cn(\neg p).$$

With p as ϕ, q as ψ, and $\neg p$ as χ, it thus follows that $(L6)$ is violated.

Unlike the last three DP-postulates, $(L4)$ to $(L6)$ seem to be fundamentally incompatible with the basic idea of infobase revision. Interestingly enough, though, infobase revision satisfies postulate $(L7)$, which is a weakened version of $(DP2)$.

PROPOSITION 7. *Every infobase revision by a satisfiable wff satisfies* $(L7)$.

We only consider the case where $\neg\phi \notin Cn(B)$. Let $\preceq$ be the $(B \circledast \neg\phi)$-induced faithful total preorder. If $\phi \notin Cn(B)$ then $B\circledast\neg\phi = B\cup\{\neg\phi\}$, and so the $(B\circledast\neg\phi)$-number (cf. Definition 3) of every model of ϕ is equal to its B-number. Therefore, $\text{Min}_{\preceq}(\phi) = M(B) \cap M(\phi) \neq \emptyset$, and so $Cn((B \circledast \neg\phi) \circledast \phi) = Cn(B) + \phi$. If $\phi \in Cn(B)$ it suffices to show that $Cn((B \circledast \neg\phi) \circledast \phi) \subseteq Cn(B)$. Pick any $x \in M(B)$. By construction, x is a model of all the wffs in $B \circledast \neg\phi$, except for $\neg\phi$, and so $x \preceq y$ for every $y \in M(\phi)$. So $M(B) \subseteq \text{Min}_{\preceq}(\phi)$, from which the required result follows.

6. Future Research

Section 3 has laid the foundation for a theory of infobase change, but it is clear that much still needs to be done. Some of the results in Section 5 make it clear that the representation of an infobase needs to be looked at again. The current representation, as a set of wffs, does not allow for the possibility of two or more instances of the same wff being regarded as independently obtained bits of information. One way of circumventing this restriction would be to represent infobases as (ordered) lists of wffs. This would also enable us to modify the definition of a B-induced faithful total preorder so that it orders interpretations according to the number of wffs in a list, and not the number of logically non-equivalent wffs, as is currently the case.

Infobase change, as we have currently defined it, assumes that the wffs contained in the base B have equal epistemic weight. But there may be good reasons for regarding some wffs in B as epistemologically more important than others, as the following example, which is part of an example in Hansson (1992b), attests to.

Example 8. “A geography student sees one of his fellow students pick up a book in the library. The title of the book is *The University at Niamey*. He asks, ‘Where is Niamey?,’ and receives the answer, ‘It is a Nigerian city.’

Next day, in an oral examination, the professor asks our student, ‘What do you know about Niamey?’ – ‘It is a university town in Nigeria’ – ‘It most certainly isn’t’ . . . the student believes what the professor says, and adjust his beliefs accordingly.”

We use the propositional language generated by the atoms p and q to represent the situation above, where p denote the assertion that there is university in Niamey, and q denotes the assertion that Niamey is a town in Nigeria. So the infobase B is $\{p, q\}$ and the student contracts $p \wedge q$ from B. It is easy to verify that an infobase contraction of B by $p \wedge q$ yields an infobase that is element-equivalent to $\{p \vee q\}$. But, as Hansson (1992b) argues, it is reasonable to assume that the result of the above contraction should be $\{p\}$. This is because of the extra-logical assumption that information obtained in library books is more reliable than information obtained from fellow students, which allows us to retain p rather than q.

One way in which to represent such extra-logical information is in terms of orderings of *epistemic relevance* on B. Nebel (1990, 1991, 1992) requires of epistemic relevance orderings to be total preorders on B. When applied to infobase change, the aim would be to use an epistemic relevance ordering on B to obtain a suitable B-faithful total preorder. An appropriate base change operation would then be constructed in a manner analogous to the way it is currently being constructed.

One of the main differences between infobase change and many other approaches to base change is illustrated by Example 2, where a wff that is not contained in the infobase $B = \{p, q\}$ finds its way into the resulting infobase $B \ominus p \wedge q$. And while this seems to be the correct solution in many respects, it is not quite in tune with the intuition that the wffs in the infobase B should represent independently obtained beliefs. For it seems counterintuitive to regard a wff that is merely entailed by the wffs in B as an independently obtained belief contained in $B \ominus p \wedge q$. It is with this kind of example in mind that Rott (1992) writes as follows (in the quotation H represents the base $\{p, q\}$):

> "... Even after conceding that one of p and q may be false, we should still cling to the belief that the other one is true. But $H' = \{p \vee q\}$ is no base which can be constructed naturally from H – it certainly does not record any explicit belief. We are faced with a deep-seated dilemma ..."

Rott ultimately decides against the inclusion of such wffs, arguing that the treatment of bases as containing explicit beliefs should get precedence.* We conclude this section by arguing that a priority ordering, similar in spirit to the epistemic relevance orderings, may provide an acceptable solution. The idea is to split the infobase obtained from an infobase contraction into a set of explicit beliefs and a set of *introspective beliefs*. After an infobase contraction of the infobase B by the wff ϕ, the set of explicit beliefs consists of the ϕ-retained wffs of B (i.e., the set $B \setminus B^{-\phi}$), while the set of introspective beliefs contains the ϕ-weakened version $w_B^{-\phi}(\psi)$ of every ϕ-discarded wff ψ of B. Wffs that were, at some stage, obtained directly from independent sources thus constitute the set of explicit beliefs, while

* Hansson (1996) mentions the use of disjunctively closed bases (in which the disjunction $\phi \vee \psi$ of every $\phi, \psi \in B$ is also in B) as a possible solution to problems of this kind. Unfortunately, this ensures that bases cannot be finite. And in any case, Hansson does not regard it as an acceptable solution, warning that it should be seen as an interesting special case, rather than a required property of bases.

wffs such as the one logically equivalent to $p \vee q$ in Example 2 are regarded as beliefs obtained by introspection during the contraction process, and are thus to be seen as carrying less epistemic weight than the explicit beliefs.

7. Conclusion

We have presented an approach to base change, termed infobase change, for which the associated theory change operations satisfy all the AGM postulates. We use the associated theory change operations to determine which wffs in the infobase should be retained after an infobase change. The wffs that cannot be retained are not simply discarded, but are weakened in an appropriate fashion. Infobase change is completely specified by an infobase and the corresponding induced faithful total preorder (which is used to obtain the associated theory change operation). In fact, since the faithful total preorder is *induced* by the infobase, all the information required for an infobase change actually resides in the infobase itself. For the purpose of performing infobase change, it is thus sufficient to represent the epistemic state of an agent by an infobase.

An advantage of the fact that an infobase generates a unique contraction and revision operation is that it allows for iterated infobase change. But this uniqueness of change operations is also a limitation since it has nothing to say about a general theory of infobase change. We contend that infobase change, as we have presented it, occupies a special place in such a general theory, perhaps analogous to that of full meet contraction in AGM theory change.

References

Alchourrón, C., Gärdenfors, P., and Makinson, D., 1985, "On the logic of theory change: Partial meet functions for contraction and revision," *Journal of Symbolic Logic* **50**, 510–530.

Borgida, A., 1985, "Language features for flexible handling of exceptions in information systems," *ACM Transactions on Database Systems* **10**, 563–603.

Boutilier, C., 1994, "Unifying default reasoning and belief revision in a modal framework," *Artificial Intelligence* **68**, 33–85.

Dalal, M., 1988, "Investigations into a theory of knowledge base revision," pp. 475–479 in *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence*, Saint Paul, MN, San Mateo, CA: Morgan Kaufmann.

Darwiche, A. and Pearl, J., 1997, "On the logic of iterated belief revision," *Artificial Intelligence* **89**, 1–29.

Fuhrmann, A., 1991, "Theory contraction through base contraction," *Journal of Philosophical Logic* **20**, 175–203.

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

Grove, A., 1988, "Two modellings for theory change," *Journal of Philosophical Logic* **17**, 157–170.

Hansson, S.O., 1989, "New operators for theory change," *Theoria* **55**, 114–132.

Hansson, S.O., 1992a, "A dyadic representation of belief," pp. 89–121 in *Belief Revision*, P. Gärdenfors, ed., Cambridge: Cambridge University Press.

Hansson, S.O., 1992b, "In defense of base contraction," *Synthese* **91**, 239–245.

Hansson, S.O., 1993, "Reversing the Levi Identity," *Journal of Philosophical Logic* **22**, 637–669.

Hansson, S.O., 1996, "Knowledge-level analysis of belief base operations," *Artificial Intelligence* **82**, 215–235.

Katsuno, H. and Mendelzon, A., 1991, "Propositional knowledge base revision and minimal change," *Artificial Intelligence* **52**, 263–294.

Lehmann, D., 1995, "Belief revision, revised," pp. 1534–1540 in *IJCAI-95, Proceedings of the 14th International Joint Conference on Artificial Intelligence*, Montreal, PQ, August 20–25, 1995, Vol. 2, C.S. Mellish, ed., San Mateo, CA: Morgan Kaufmann.

Nayak, A.C., 1994, "Foundational belief change," *Journal of Philosophical Logic* **23**, 495–533.

Nebel, B., 1989, "A knowledge level analysis of belief revision," pp. 301–311 in *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, R.J. Brachman, H.J. Levesque, and R. Reiter, eds., San Mateo, CA: Morgan Kaufmann.

Nebel, B., 1990, *Reasoning and Revision in Hybrid Representation Systems*, Lecture Notes in Artificial Intelligence, Vol. 422, Berlin: Springer-Verlag.

Nebel, B., 1991, "Belief revision and default reasoning: Syntax-based approaches," pp. 417–428 in *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference KR '91)*, J. Allen, R. Fikes, and E. Sandewall, eds., San Francisco, CA: Morgan Kaufmann.

Nebel, B., 1992, "Syntax-based approaches to belief revision," pp. 52–88 in *Belief Revision*, P. Gärdenfors, ed., Cambridge Tracts in Theoretical Computer Science, Vol. 29, Cambridge: Cambridge University Press.

Rott, H., 1992, "Modellings for belief change: Prioritization and entrenchment," *Theoria* **58**, 21–57.

Satoh, K., 1988, "Nonmonotonic reasoning by minimal belief revision," pp. 455–462 in *Proceedings International Conference on Fifth Generation Computer Systems*, Tokyo.

Weber, A., 1986, "Updating propositional formulas," pp. 487–500 in *Proceedings First Conference on Expert Database Systems.*

Winslett, M., 1988, "Reasoning about action using a possible models approach," pp. 89–93 in *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence*, Saint Paul, MN, Palo Alto, CA: Morgan Kaufmann.