

Considerations on a Similarity-Based Approach to Belief Change

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

Abstract

A foundational approach to modelling belief contraction and revision is presented, based on a notion of *similarity* between belief sets. In contracting α from a belief set, the result is the belief set(s) most similar to the original in which α is not believed; similar considerations apply to belief revision. The modelling of belief change generalises the Grove modelling based on a system of spheres, where instead of having a total order on sets of possible worlds, we have a total order on sets of belief sets. Given this modelling, sets of postulates are determined for contraction and revision. The resulting postulate sets subsume those in the AGM approach. The approach sheds light on the foundations of belief revision in that, first, it provides a more general framework than the AGM approach; second, it illustrates assumptions underlying the AGM approach; and third, it allows a “fine-grained” investigation of proposed principles underlying belief change. Lastly, it demonstrates that, at their most basic, revision and contraction of beliefs are not interdefinable notions, but rather distinct concepts.

1 Introduction

A belief set of facts, assertions, etc. will of course change over time with the addition or deletion of information. In this paper I am concerned with a foundational characterisation of belief contraction and revision. The questions addressed are familiar: given a belief set and a sentence to be added to the belief set, what can we say about the revised belief set? And: given a belief set and a sentence to be contracted, again, what can be said about the result? For reasons that will become apparent, I focus on belief contraction rather than the revision of beliefs.

Two assumptions are made in addressing these questions. First, that a change is *successful*, so that after a sentence is contracted from a belief set, that sentence is no longer believed in the resulting belief set(s). Second, I assume that belief change is founded on a notion of *similarity* among belief sets. Thus if a sentence α is

to be contracted from (added to) belief set K , then the result will be the belief set most similar to K in which α is not believed (is believed). This reduces the notion of belief change to that of similarity between belief sets. The focus here is on an abstract characterisation of beliefs, so we will be concerned with syntax-independent characterisations of change functions.

Clearly, there is not a great deal that can be said in general about similarity between belief sets. In the approach presented here, it is assumed that every belief set K has associated with it a binary metric of relative similarity to K , and that this metric is a total preorder. That is, if $\preceq_K$ is the relative similarity metric associated with K then $\preceq_K$ is reflexive, transitive and connected; as well K is the minimum element in the order. I assume further that for every sentence α there is a $\preceq_K$ -least belief set or set of belief sets in which α is believed. Given such a similarity order, contraction and revision functions can be defined. Subsequent to this modelling, corresponding postulates validated by these functions are determined.

The approach is intended to provide a *minimal* notion of belief change, in that for the postulates that obtain, all should arguably hold for any syntax-independent, successful change functions. Consequently, the approach is more basic than, and so subsumes, the AGM approach to belief change. However, as shown at the end of Section 3, the approach does not subsume the Katsuno-Mendelzon approach to belief update. A benefit of the present approach is that it allows a very “fine-grained” investigation of principles underlying belief revision where distinct notions are, in fact, distinguished. Consequently the semantics illustrates that in the AGM approach, there are a number of distinct (albeit very basic, and perhaps beyond debate) principles composing the approach. Secondly, the approach is arguably intuitive and plausible, in that it is based on commonsense intuitions regarding belief revision and contraction. By imposing constraints on the semantic theory, additional postulates may be satisfied. Arguably, such constraints will reflect plausible intuitions concerning belief change, and so the approach will help provide insight into different belief change functions. Finally, and perhaps surprisingly, at this very basic level it proves to be the case that revision and contraction functions comprise distinct notions, with

contraction being the more general. This is in contrast with the AGM approach, where revision and contraction are in a certain sense interdefinable.

One omission in this paper is that iterated revision is not addressed. The reason for this is that, at this point, our interests lie with a comparison to the AGM approach and, for the present, uniterated revision. The final section briefly considers how the notion of similarity may be used in iterated revision.

Section 2 briefly reviews the AGM approach and the Grove construction. Section 3 presents the approach, while Section 4 provides a conclusion.

2 Background

Belief sets change over time, with the addition and deletion of information. In general, there is no purely logical reason for making one choice rather than another among the sentences to be retracted or kept. Hence from a logical view there may be several ways of specifying a belief change function. However, general properties of such functions can be investigated.

In the *AGM approach* of Alchourron, Gärdenfors, and Makinson [AGM85; Gär88], standards for revision and contraction functions are given by various *rationality postulates*. The goal is to describe belief change at the *knowledge level*, that is on an abstract level, independent of how beliefs are represented and manipulated. Belief states are modelled by sets of sentences closed under the logical consequence operator of some logic in some language L , where the logic includes classical propositional logic. A *belief set* is a set K of sentences which satisfies the constraint: If K logically entails β then $\beta \in K$. $K + \alpha$ is the deductive closure of $K \cup \{\alpha\}$, and is called the *expansion* of K by α . $K_{\perp}$ is the inconsistent knowledge base (i.e. $K_{\perp} = L$). $\mathcal{T}$ is the set of all belief sets.

For *contraction*, some beliefs are retracted but no new beliefs are added. A contraction function $\dot{-}$ is a function from $\mathcal{T} \times L$ to $\mathcal{T}$ satisfying the following postulates.

- (K $\dot{-}$ 1) $K \dot{-} \phi$ is a belief set.
- (K $\dot{-}$ 2) $K \dot{-} \phi \subseteq K$.
- (K $\dot{-}$ 3) If $\phi \notin K$, then $K \dot{-} \phi = K$.
- (K $\dot{-}$ 4) If not $\vdash \phi$, then $\phi \notin K \dot{-} \phi$.
- (K $\dot{-}$ 5) If $\phi \in K$, then $K \subseteq (K \dot{-} \phi) + \phi$.
- (K $\dot{-}$ 6) If $\vdash \phi \equiv \psi$, then $K \dot{-} \phi = K \dot{-} \psi$.
- (K $\dot{-}$ 7) $K \dot{-} \phi \cap K \dot{-} \psi \subseteq K \dot{-} (\phi \wedge \psi)$.
- (K $\dot{-}$ 8) If $\psi \notin K \dot{-} (\phi \wedge \psi)$, then $K \dot{-} (\phi \wedge \psi) \subseteq K \dot{-} \psi$.

A *revision* function $\dot{+}$ is a function from $\mathcal{T} \times L$ to $\mathcal{T}$ 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 $(K \dot{+} \alpha) + \beta \subseteq K \dot{+} (\alpha \wedge \beta)$.

Katsuno and Mendelzon [KM92] explore a distinct notion of belief change, comprising belief *update* and *erasure*, wherein an agent changes its beliefs in response to changes in its external environment. Our interests here centre on the AGM approach; however in Section 3.4, I briefly consider this approach.

In [Gro88] a modelling of the AGM postulates is given based on Lewis' *system of spheres* semantics [Lew73]. M_L is the set of all maximal consistent sets of sentences of L . Intuitively, an element of M_L can be thought of as corresponding to an interpretation in the language, or alternatively to a possible world. Define $|\alpha| = \{I \in M_L \mid \alpha \in I\}$ for system of spheres M .

Definition 2.1 ([Gro88]) A set of subsets $\mathcal{S}$ of M_L is a system of spheres centred on X where $X \subseteq M_L$, if it satisfies the conditions:

- S1 $\mathcal{S}$ is totally ordered by $\subseteq$.
- S2 X is the minimum of $\mathcal{S}$ under $\subseteq$.
- S3 $M_L \in \mathcal{S}$.
- S4 If $|\alpha| \neq \emptyset$ then there is a least (wrt $\subseteq$) sphere $c(\alpha)$ such that $c(\alpha) \cap |\alpha| \neq \emptyset$ and $J \cap |\alpha| \neq \emptyset$ implies $c(\alpha) \subseteq J$ for every $J \in \mathcal{S}$.

$f_S(\alpha)$ is defined to pick out the least (if such there be) interpretations containing α ; i.e. $f_S(\alpha) = |\alpha| \cap c(\alpha)$. The principal result is a correspondence between systems of spheres and the AGM postulates, in that, informally, for any system of spheres centred on $|K|$ there is a corresponding revision function that satisfies the AGM postulates and, conversely, for any revision function satisfying the AGM postulates there is a corresponding system of spheres centred on $|K|$.

In the next section, we take the Grove modelling as our point of departure, essentially advocating a modelling based on a system of spheres but where belief sets replace possible worlds in the modelling.

3 Similarity Orderings on Belief Sets

Contraction of belief sets is addressed first, followed by *revision*. The central intuition is that in contracting α from K , we want to select the *most similar* belief set(s) to K in which α is not believed. (As a point of interest, the AGM approach assumes that one wants to retain as much of the information in the belief set as possible; this criterion would constitute a specific similarity measure.) A belief set K has associated with it a (binary) metric $\preceq_K$ of relative similarity to K , and this metric is a total preorder. As well K is the minimum element in the order. For every α there is a $\preceq_K$ -least belief set or set of belief sets in which α is believed. This last constraint is analogous to the Limit Assumption of [Lew73].

For contracting α from K , we select the belief sets most similar to K in which α is not believed. Since there

may be more than one such belief set, and since there is nothing to distinguish these belief sets, the contraction of α from K corresponds to this set of belief sets. This is in contrast with the AGM approach, where a contraction function has as value a single belief set.

3.1 A Modelling for Belief Change

$K, K_1, K_2, \dots$ will denote belief sets. Recall that $\mathcal{T}$ is the set of belief sets. When we come to the revision and contraction postulates, it will be convenient to be able to talk about the belief sets in which a sentence α is believed.

Definition 3.1 $KB(\alpha) = \{K \mid K \in \mathcal{T} \text{ and } \alpha \in K\}$

Definition 3.2 A similarity order model on belief sets centred on K is given by $M = \langle K, \mathcal{KB}, \preceq_K \rangle$ where

1. $K \in \mathcal{KB}$
2. $\mathcal{KB} \subseteq \mathcal{T}$ and $K_1 \in \mathcal{KB}$
3. $\preceq_K \subseteq (\mathcal{KB} \times \mathcal{KB})$ where $\preceq_K$ satisfies the conditions:
 - P1** $\preceq_K$ is connected (and so reflexive). So for $K_1, K_2 \in \mathcal{KB}$, we have $K_1 \preceq_K K_2$ or $K_2 \preceq_K K_1$.
 - P2** $\preceq_K$ is transitive. That is, if $K_1 \preceq_K K_2$ and $K_2 \preceq_K K_3$ then $K_1 \preceq_K K_3$.
 - P3** K is the minimum of $\preceq_K$. So for every $K_1 \in \mathcal{KB}$, $K \preceq_K K_1$, and if $K_1 \preceq_K K$ then $K_1 = K$.
 - P4** If $\nabla \alpha$ then there is $K_1 \in \mathcal{KB}$ such that $K_1 \nabla \alpha$ and for every K_2 if $K_2 \nabla \alpha$ then $K_1 \preceq_K K_2$.

P4 states that for every sentence α there is a $\preceq_K$ -least belief set or belief sets in which α is believed. This condition is analogous to $(S \leq 3)$ of [Gro88], expressing the Limit Assumption.

3.2 Belief Change: Contraction

$\min C(\alpha)$ is defined as the least set of belief sets in which α is consistent.

Definition 3.3 $\min C(\alpha) = \{K_1 \in \mathcal{T} \mid K_1 \nabla \alpha \text{ and if } K_2 \nabla \alpha \text{ then } K_1 \preceq_K K_2\}$.

We can now define belief contraction.

Definition 3.4 The contraction of α from theory K in M is given by:

$$K \dot{\alpha} = \begin{cases} \min C(\neg\alpha) & \text{if } \nabla \alpha \\ K & \text{if } \vdash \alpha \end{cases}$$

Given this semantics we can ask what postulates are satisfied. For reference I distinguish the postulates in a definition. The numbering is with reference to the corresponding (or most similar) AGM postulates.

Definition 3.5 The following constitute the set of KB contraction postulates.

- $(K \dot{\alpha} 1)$ $K \dot{\alpha}$ is a non-empty set of belief sets.
- $(K \dot{\alpha} 3)$ If $K \nabla \alpha$ then $K \dot{\alpha} = \{K\}$.
- $(K \dot{\alpha} 4)$ If $\nabla \alpha$ then $K \dot{\alpha} \cap KB(\alpha) = \emptyset$.
- $(K \dot{\alpha} 6)$ If $\vdash \alpha \equiv \beta$ then $K \dot{\alpha} = K \dot{\beta}$.

$(K \dot{\alpha} 7)$ If $\nabla \neg(\alpha \wedge \beta)$ then: if $K \dot{\beta} \not\subseteq KB(\alpha)$ then $K \dot{\alpha} = K \dot{(\alpha \wedge \beta)} \setminus KB(\alpha)$.

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

Postulates $(K \dot{\alpha} 3)$, $(K \dot{\alpha} 4)$, and $(K \dot{\alpha} 6)$ are essentially the same as their AGM counterparts. For $(K \dot{\alpha} 1)$, contraction isn't guaranteed to result in a single belief set. This corresponds to the fact that in the semantics there may be minimal, equivalently-similar belief sets in which α isn't believed. The AGM postulate $(K \dot{\alpha} 2)$ reflects the requirement that no new beliefs occur in a contraction. In our case, the most similar belief set(s) to K in which α isn't believed may indeed contain new information. Consider for example a nonmonotonic belief set wherein $Bird(Opus)$, $Penguin(Opus)$, and $\neg Fly(Opus)$ are believed. If $Penguin(Opus)$ is contracted, then if we have the usual default rules concerning birds and flying, we might elect to replace $\neg Fly(Opus)$ by $Fly(Opus)$ in the resultant belief set. $(K \dot{\alpha} 3)$ reflects the fact that a belief set is most similar to itself, and so if $\neg\alpha$ is consistent with K , contracting α results in the set of K . $(K \dot{\alpha} 4)$ asserts that contraction is successful while $(K \dot{\alpha} 6)$ reflects the fact that it is the content of α that determines the contraction and not its syntactic expression.

The AGM recovery postulate $(K \dot{\alpha} 5)$ is missing: if $K \vdash \alpha$ then K and $K \dot{\alpha}$ may be quite different and, in fact, $K \dot{\alpha}$ may contain information not contained in K (since we don't have an equivalent to $(K \dot{\alpha} 2)$). Hence $(K \dot{\alpha} \alpha) + \alpha$ may be different from K . It proves to be the case that equivalents to the AGM postulates $(K \dot{\alpha} 7)$ and $(K \dot{\alpha} 8)$ are consequences of the KB contraction postulates. The quite different-looking $(K \dot{\alpha} 7)$ and $(K \dot{\alpha} 8)$ are employed because they readily yield a representation result. $(K \dot{\alpha} 7)$ is quite strong: if the contraction of β from K results in at least one belief set in which $\neg\alpha$ is consistent, then $K \dot{\alpha}$ consists of just those belief sets in $K \dot{(\alpha \wedge \beta)}$ in which $\neg\alpha$ is consistent. Semantically, $(K \dot{\alpha} 8)$ asserts that if $K \dot{\alpha}$ and $K \dot{(\alpha \wedge \beta \wedge \gamma)}$ are equivalently-similar then so is $K \dot{(\alpha \wedge \beta)}$.

There are number of interesting results following from these postulates. We have already noted that $(K \dot{\alpha} 7)$ and $(K \dot{\alpha} 8)$ are consequences of the postulates. A selection of other results are given in the following theorem.

Theorem 3.1

1. If $\nabla \alpha$ then: $K \dot{\alpha} = K \dot{\beta}$ iff $K \dot{\alpha} \cap KB(\beta) = \emptyset$ and $K \dot{\beta} \cap KB(\alpha) = \emptyset$.
2. $K \dot{(\alpha \wedge \beta)} = K \dot{\alpha} \cup K \dot{\beta}$ or $K \dot{(\alpha \wedge \beta)} = K \dot{\alpha}$ or $K \dot{(\alpha \wedge \beta)} = K \dot{\beta}$.
3. If $K \dot{\alpha} \subseteq K \dot{\beta}$ then $K \dot{(\alpha \wedge \gamma)} \subseteq K \dot{(\beta \wedge \gamma)}$.
4. If $K \dot{\alpha} \subseteq K \dot{(\alpha \wedge \beta)}$ and $K \dot{\beta} \subseteq K \dot{(\beta \wedge \gamma)}$ then $K \dot{\alpha} \subseteq K \dot{(\alpha \wedge \gamma)}$.

The first result states that if every belief set in $K \dot{\alpha}$ is consistent with $\neg\beta$, and if every belief set in $K \dot{\beta}$ is consistent with $\neg\alpha$, then $K \dot{\alpha}$ and $K \dot{\beta}$ coincide. The second result is analogous to the "factoring" result of the AGM postulates [Gär88, p. 57 (3.27)], but expressed in

terms of sets of belief sets. The final result essentially expresses a notion of transitivity in our representation theorem; this in turn is based on the fact that if $K \dot{\alpha} \subseteq K \dot{\alpha} \wedge \beta$ then any minimal belief set in which $\neg\alpha$ is consistent is also a minimal belief set in which $\neg(\alpha \wedge \beta)$ is consistent, justifying an assertion that $K \dot{\alpha} \preceq_K K \dot{\alpha} \wedge \beta$.

We obtain the following results relating the KB contraction postulates to similarity order models.

Theorem 3.2 *Let M be any similarity order model on belief sets centred on K . If $K \dot{\alpha}$ is defined according to Definition 3.4 then the KB contraction postulates are satisfied in M .*

Theorem 3.3 *Let $\dot{\alpha}$ be a function from $\mathcal{T} \times L$ to $2^{\mathcal{T}}$ satisfying the KB contraction postulates. Then for any fixed theory $K \in \mathcal{T}$ there is a similarity order model on belief sets centred on K satisfying Definition 3.4 for all $\alpha \in L$.*

We can determine what conditions are required to recover the other AGM postulates. This can be accomplished in two ways. First, we can consider criteria which satisfy individual postulates. Second, we can consider a criterion that would *en masse* as it were, yield the AGM postulates. (A third, and most interesting, possibility is given in Theorem 3.8 in the final subsection of this section.)

In the first case, for example, one can obtain a postulate equivalent to $(K \dot{-} 2)$ by restricting the similarity relation to belief sets strictly weaker than K . To obtain a postulate equivalent to $(K \dot{-} 1)$ there are various strategies that can be employed. To obtain a single belief set from a contraction, one could define some *selection function* that returns a single belief set given a set of equally-similar belief sets. For example, this function could select an arbitrary belief set, or it might select a belief set on the basis of some other criterion, for example, the overall *simplicity* of the belief set. Or it might determine some representative belief set, for example, the intersection of the set.

Alternately, one might decide that the semantics be refined so that the contraction function returns a single belief set. Again, there are various alternatives. For example, one could require that the similarity order be antisymmetric, so that if $K_i \preceq_K K_j$ then $K_j \not\preceq_K K_i$ unless $K_i = K_j$. Alternately, one could require that equally-similar belief sets be closed under intersection, so that if K_i and K_j were equally-similar to K then so would be $K_i \cap K_j$; the contraction then would return the minimal (in terms of containment) belief set. For either strategy, the imposition of additional constraints would not be ad hoc, but rather should reflect reasonable assumptions in the semantics. So if $\preceq_K$ were a total order on belief sets, then one would be compelled to accept the assumption that there are no “ties” in similarity of belief sets. If one decided that contraction is closed under intersections, then presumably one should be able to justify this choice. The point here is that the approach allows such distinctions to be made.

On the other hand, one obtains the full set of AGM contraction postulates by asserting that successive weakenings of K are less similar to K . This reflects a criterion of *informational economy*, that we retain as much as possible of our old beliefs. The use of a selection function, below, is one of a number of ways to guarantee that a single belief set results from a contraction. We obtain:

Theorem 3.4 *Let $M = \langle K, \mathcal{KB}, \preceq_K \rangle$ be a similarity order model on belief sets centred on K such that: $\mathcal{KB} = \{K' \in \mathcal{T} \mid K' \subseteq K\}$ and if $K' \subseteq K''$ then $K'' \preceq_K K'$. Let S be any selection function on $K \dot{\alpha}$ such that $S(K \dot{\alpha}) = \cap S'$ for some $S' \subseteq K \dot{\alpha}$. Then the AGM contraction postulates are satisfied in M by the function $S(K \dot{\alpha})$.*

Note that in the above theorem, if we allowed an arbitrary similarity order model, and we defined a new contraction operator to be $\bigcap_{K' \in (K \dot{\alpha})} K'$ (as suggested for example in [Neb92]) that the only new AGM postulate satisfied is $(K \dot{-} 1)$.

3.3 Belief Change: Revision

We turn now to belief revision. The main result is that, surprisingly, this function is not interdefinable with contraction, and in fact is weaker.

Given a similarity order model M , we define $\min(\alpha)$ as the least set of belief sets in which α is true, following which we define belief revision.

Definition 3.6 $\min(\alpha) = \{K_1 \in \mathcal{KB}(\alpha) \mid K_1 \preceq_K K_2 \text{ for every } K_2 \in \mathcal{KB}(\alpha)\}$.

Definition 3.7 *The revision of theory K by α in M is given by: $K \dot{\alpha} = \min(\alpha)$.*

Revision postulates are given in the next definition, with numbering in reference to the corresponding (or most similar) AGM revision postulates.

Definition 3.8 *The following constitute the set of KB revision postulates.*

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

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

$(K \dot{\alpha} 3)$ If $K \vdash \alpha$ then $K \dot{\alpha} = \{K\}$.

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

$(K \dot{\alpha} 7)$ If $K \dot{\alpha} \cap \mathcal{KB}(\beta) \neq \emptyset$ then $K \dot{\alpha} \wedge \beta = K \dot{\alpha} \cap \mathcal{KB}(\beta)$.

$(K \dot{\alpha} 8)$ If $K \dot{\alpha}_i \cap \mathcal{KB}(\alpha_{i+1}) \neq \emptyset$ for $0 \leq i < n$, and $K \dot{\alpha}_n \cap \mathcal{KB}(\alpha_0) \neq \emptyset$ then $K \dot{\alpha}(\alpha_0 \wedge \alpha_n) \subseteq K \dot{\alpha}_0$.

$(K \dot{\alpha} 2)$ and $(K \dot{\alpha} 6)$ are the only postulates the same as their AGM counterparts. For $(K \dot{\alpha} 1)$, revision, like contraction, isn't guaranteed to result in a unique belief set. $(K \dot{\alpha} 2)$ reflects the requirement that the revision be successful. $(K \dot{\alpha} 3)$ is an obvious weakening of $(K \dot{+} 3)$; it is difficult to think of a situation where it shouldn't hold. In contrast, it seems feasible that a revision function may not satisfy the AGM postulates $(K \dot{+} 3)$ and $(K \dot{+} 4)$, since if α is consistent with K , it may be that $K + \alpha$ isn't the most similar belief set to K in which α is believed; for

this to hold we could again bring in an assumption of informational economy. $(K+5)$ is missing: if $\not\models \neg\alpha$ there is nothing forbidding $K \dot{\circ} \alpha = K_{\perp}$. $(K \dot{\circ} 7)$ and $(K \dot{\circ} 8)$ again are dissimilar from their AGM counterparts.

Again, various reasonable and interesting results following from these postulates. Several examples are given in the following theorem.

Theorem 3.5

1. $K \dot{\circ} \alpha = K \dot{\circ} \beta$ iff $K \dot{\circ} \alpha \subseteq KB(\beta)$ and $K \dot{\circ} \beta \subseteq KB(\alpha)$.
2. If $K \dot{\circ} \alpha \subseteq KB(\beta)$ then $K \dot{\circ} (\alpha \wedge \beta) = K \dot{\circ} \alpha$.
3. If $K \dot{\circ} \alpha_i \subseteq KB(\alpha_{i+1})$ for $0 \leq i < n$ and $K \dot{\circ} \alpha_n \subseteq KB(\alpha_0)$ then $K \dot{\circ} \alpha_i = K \dot{\circ} \alpha_{i+1}$ for $0 \leq i < n$.
4. $K \dot{\circ} \alpha \cap K \dot{\circ} \beta = \emptyset$ or $K \dot{\circ} (\alpha \wedge \beta) \subseteq K \dot{\circ} \alpha$ or $K \dot{\circ} (\alpha \wedge \beta) \subseteq K \dot{\circ} \beta$.

For the third result, if $K \dot{\circ} \alpha_i \subseteq KB(\alpha_{i+1})$ then semantically (see below) the least α_i belief sets are no more similar to K than the least α_{i+1} belief sets. If a chain of such containments forms a “loop”, then the revisions are equally similar and, in fact, equal. The final result provides a weaker version of the AGM “factoring” result (see [Gär88, p. 57 (3.16)]). It is also weaker than the corresponding result for contraction (Theorem 3.1.2). As well, the AGM postulates $(K+7)$ and $(K+8)$ are not logical consequences of the KB revision postulates, in contradistinction to the AGM contraction postulates $(K-7)$ and $(K-8)$ whose analogues are logical consequences of the KB contraction postulates.

In addition, we do not obtain the representation result for revision that we do for contraction. Define a *weak similarity order model on belief sets centred on K* to be a similarity order model on belief sets *except* rather than being connected, it is reflexive only. We obtain the following results relating the KB revision postulates to weak similarity order models.

Theorem 3.6 *Let M be a weak similarity order model on belief sets centred on K . If we define $K \dot{\circ} \alpha = \min(\alpha)$ then the KB revision postulates are satisfied in M .*

Theorem 3.7 *Let $\dot{\circ}$ be a function from $\mathcal{T} \times \mathcal{L}$ to $2^{\mathcal{T}}$ satisfying the KB revision postulates. Then for any fixed theory $K \in \mathcal{T}$ there is a weak similarity order model on belief sets centred on K satisfying $K \dot{\circ} \alpha = \min(\alpha)$ for all $\alpha \in \mathcal{L}$.*

As with contraction we can ask what conditions are required to recover the other AGM postulates. We can specify that revision has a unique belief set as its value via strategies sketched previously. To obtain an equivalent of $(K+5)$ we would require (and not unreasonably) that the inconsistent belief set be the most dissimilar of belief sets to any consistent belief set. Other postulates are dealt with by imposing similar conditions.

3.4 Discussion

As in the AGM approach, the present approach leaves open how a specific contraction or revision function may

be defined. Rather, the approach provides constraints that contraction and revision functions must obey. If one accepts that a notion of similarity as developed here underlies belief change then one presumably would accept the respective postulate sets that would limit properties of an acceptable change function. One can further restrict the class of acceptable functions by placing additional restrictions on the notion of similarity. Thus if the range of a contraction function for belief set K is restricted to be a subsumed belief set of K , this together with a selection function restricts the satisfying contraction functions to those satisfying the AGM postulates.

On the other hand, one could propose a specific metric of similarity for a (say) revision operator. For example, if we equated a belief set with a set of possible worlds rather than a set of sentences, then Dalal’s approach [Dal88] is easily expressed using similarity: for belief set K the most similar belief sets to K not the same as K would be those composed of possible worlds differing in one literal from a world in K . The next closest set of belief sets would be those composed of worlds differing in two literals from a world in K , and so on. The result of revising K by α would be the maximal, nearest belief set in which α is true.

As mentioned, revision proves to be weaker than contraction. In detail, in the proof of Theorem 3.3 a similarity order model is defined such that for belief sets K_1 and K_2 there are α and β such that $K_1 \in K \dot{\circ} \alpha$ and $K_2 \in K \dot{\circ} \beta$. Theorem 3.1.2 yields that for every α and β : $K \dot{\circ} \alpha \subseteq K \dot{\circ} (\alpha \wedge \beta)$ or $K \dot{\circ} \beta \subseteq K \dot{\circ} (\alpha \wedge \beta)$. $K \dot{\circ} \alpha \subseteq K \dot{\circ} (\alpha \wedge \beta)$ for example asserts that every belief set in $K \dot{\circ} \alpha$ is in $K \dot{\circ} (\alpha \wedge \beta)$, or informally the belief sets in $K \dot{\circ} \alpha$ and $K \dot{\circ} (\alpha \wedge \beta)$ are equivalently similar, so the belief sets in $K \dot{\circ} \beta$ are no more similar to K than the belief sets in $K \dot{\circ} (\alpha \wedge \beta)$, whence $K_1 \preceq_K K_2$. For revision, there is no such relation among belief sets, and the belief sets in $K \dot{\circ} \alpha$, $K \dot{\circ} \beta$, $K \dot{\circ} (\alpha \wedge \beta)$, and $K \dot{\circ} (\alpha \vee \beta)$ may be distinct. We thus lose the capability to define connectivity in Theorem 3.7. This difference in turn relies on the fact that contraction yields belief sets *consistent* with a sentence, whereas revision yields belief sets in which a sentence is *provable*.

Interestingly, a revision operator satisfying the AGM postulates is obtained in terms of KB contraction and the Levi identity, but AGM contraction is not recoverable from KB revision using the Harper identity.

Theorem 3.8 *Let M be any similarity order model centred on K with $K \dot{\circ} \alpha$ defined as in Definition 3.4 where*

1. $\dot{\circ} : \mathcal{KB} \times \mathcal{L} \mapsto \mathcal{KB}$
2. $K_1 \preceq_K K_{\perp}$ for every $K_1 \in \mathcal{KB}$.

Define $K \dot{\circ} \alpha = (K \dot{\circ} \neg\alpha) + \alpha$. Then $K \dot{\circ} \alpha$ satisfies the AGM revision postulates.

See the discussion at the end of Section 3.2 for meeting the first proviso in the theorem; the second proviso states that the incoherent belief set is maximally dissimilar to K . The theorem is interesting, in that it arguably demonstrates that AGM revision is founded on

assumptions of similarity (as given in KB contraction) *plus* informational economy (as implicit in the expansion in the Levi identity) *plus* uniqueness (proviso one) *plus* the avoidance of incoherent belief states (proviso two).

We don't obtain a similar result for revision and the Harper identity.

Theorem 3.9 *Let M be any similarity order model centred on K with $K \dot{\alpha}$ defined as in Definition 3.7 where*

1. $\dot{\alpha} : KB \times L \mapsto KB$
2. $K_1 \preceq_K K_\perp$ for every $K_1 \in KB$.

Define $K \dot{\alpha} = K \cap (K \dot{\alpha} \neg \alpha)$. Then $K \dot{\alpha}$ satisfies $(K \dot{-} 1)$, $(K \dot{-} 2)$, $(K \dot{-} 4)$ provided $K \neq K_\perp$, and $(K \dot{-} 6)$.

Not surprisingly, the KB contraction/revision functions are weaker than the corresponding AGM functions.

Theorem 3.10 *Let f be a function satisfying the AGM contraction (revision) postulates. Then f satisfies the KB contraction (revision) postulates.*

More surprising is the fact that the contraction postulates are not strictly weaker than the KM erasure postulates [KM92], in that $(K \dot{-} 3)$ is not a consequence of the erasure postulates. However the KB revision postulates subsume the update postulates.

Theorem 3.11 *Let f be a function satisfying the KM update postulates. Then f satisfies the KB revision postulates.*

It is an interesting, but unexplored, question to determine whether there is anything about contraction and revision as defined here that lends them most naturally to the AGM and KM approaches respectively.

4 Conclusion

A foundational approach has been presented in which to investigate belief change. The central intuition is that change to a belief set K by a sentence α is with reference to the belief set(s) most similar to K . The approach is quite basic, in that various of the AGM postulates don't hold, or only hold in a weaker form. Arguably the approach is not too basic, in that interesting properties still obtain, as given in the set of KB contraction and revision postulates. Moreover, the approach allows fine-grained control over the properties of contraction and revision functions. This is illustrated by the fact that of the basic AGM postulates that don't hold in the approach, each may be independently satisfied in some augmentation of the approach. An advantage of the approach then, as a foundational approach to revision, is that while the semantic basis is intuitive, such additional assumptions must be explicitly recognised and made. As a corollary, the approach arguably demonstrates that AGM revision can be viewed as being founded on a number of distinct assumptions including similarity, informational economy, and the avoidance of incoherent belief states.

A further result of this inquiry is that it appears that, at their core, revision and contraction constitute distinct

functions, with revision being the weaker. A question for future work to ask what it is about the AGM approach that leaves revision and contraction there interdefinable but not here. A second question concerns the relation of the approach to update and erasure.

There has been substantial recent interest in iterated belief revision. Iterated revision has not been addressed here, mainly because our foremost interest is in developing an approach that in some sense is more basic than the AGM approach. Clearly iteration could be addressed by investigating relations among similarity orders; in fact it may be that iterated change is more easily addressed here than in the AGM approach, primarily because here we have stepped back from some of the commitments of the AGM approach. A straightforward approach for incorporating iterated revision is, in a model, to define a mapping from pairs of belief sets to ordinals, giving the relative similarity of every pair of belief sets. From this it is an easy step to define an *epistemic state* for a belief set, K , corresponding to the total preorder expressing the relative similarity of each belief set to K .

Acknowledgements

I thank Maurice Pagnucco for his extensive and very helpful comments on an earlier version of this paper. As well I thank two reviewers for their comments.

References

- [AGM85] C.E. Alchourron, 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.
- [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.
- [Gär88] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.
- [Gro88] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [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.
- [Lew73] D. Lewis. *Counterfactuals*. Harvard University Press, 1973.
- [Neb92] B. Nebel. Syntax based approaches to belief revision. In P. Gärdenfors, editor, *Belief Revision*, pages 52–88. Cambridge University Press, 1992.