

Shifting Priorities: Simple Representations for Twenty-seven Iterated Theory Change Operators

Abstract. Prioritized bases, i.e., weakly ordered sets of sentences, have been used for specifying an agent's 'basic' or 'explicit' beliefs, or alternatively for compactly encoding an agent's belief state without the claim that the elements of a base are in any sense basic. This paper focuses on the second interpretation and shows how a shifting of priorities in prioritized bases can be used for a simple, constructive and intuitive way of representing a large variety of methods for the change of belief states — methods that have usually been characterized semantically by a system-of-spheres modeling. Among the methods represented are 'radical', 'conservative' and 'moderate' revision, 'revision by comparison' in its raising and lowering variants, as well as various constructions for belief expansion and contraction. Importantly, none of these methods makes any use of numbers.

Keywords: theory change, belief bases, belief revision, prioritization, iteration, two-dimensional revision operators.

1. Introduction

“All necessary reasoning without exception is diagrammatic,” said Charles Sanders Peirce (1903, p. 212). According to Peirce, the only way of understanding logical and mathematical propositions is by perceiving generalities in diagrams. The history of belief revision seems to confirm this thesis. By far the most intuitive representation of what is involved in various operations of belief change uses a modelling by means of systems of spheres (briefly, *SOS*) in the style of Lewis (1973) and Grove (1988). The *SOS* picture, however, is not without disadvantages. First, while it is excellently suited for the representation of the *changes* of belief states, it does not make for an easy grasp of the *contents* of the belief states in question. Second, *SOS*'s are sets of sets of large cardinalities. A more constructive approach would seem to be welcome in order to turn our semantic intuitions into something more manageable. Third, it is not evident at all where the systems of spheres of possible worlds come from.

Prioritized (or 'stratified') bases, on the other hand, have been used (i) for the representation of an agent's explicit beliefs (e.g., in Rescher 1964,

Nebel 1992, Rott 1992, Dubois, Lang and Prade 1994, Williams 1995) as well as (ii) for the compact encoding of belief states (e.g., in Rott 1991b). The motivating ideas are quite different in the two cases. In interpretation (i), it makes an essential difference whether one has p and q separately or conjoined into $p \wedge q$ in the belief base, in interpretation (ii) these are just notational variants without a difference in “meaning”. Still the most elaborate account of the first interpretation of belief bases (without prioritization) is due to Hansson (1999). In this paper we are only interested in the second interpretation. Prioritized bases have been used to represent *single* belief states. In this paper, I will explain how they can be used in what appears to me a very elegant way of representing a large variety of *changes of* belief states.

Once one has a syntactic representation that corresponds to the semantic SOS modelling of single belief states, it is natural to ask whether there are operations on these syntactic representations that correspond to reasonable transformations of SOS’s. This is the topic of this paper.

2. Representing doxastic states: Prioritized belief bases, entrenchment, systems of spheres

A *prioritized belief base* is a sequence of sets of sentences $\vec{H} = \langle H_1, \dots, H_n \rangle$. For $i < j$, the elements in H_j are supposed to be more “certain” or “reliable” or more “important” than the elements in H_i , while the elements within each H_i are tied. We presume that there are no incomparabilities. We shall also frequently use the alternative notation

$$\vec{H} = H_1 \prec \dots \prec H_n$$

This generates, in an obvious way, a transitive and complete ordering \preceq between the H_i ’s, and also between the elements of the H_i ’s.

If \vec{H} were intended to be a belief base representing the *explicit beliefs* of an agent, then the syntactical structure of the elements in each H_i would be important. In this paper, however, I am only interested in prioritized belief bases as compact and convenient *representations* of doxastic states (in structured *axiomatizations* as it were). Let us assume in the following for the sake of simplicity that not only the number of H_i ’s but also each of the individual sets H_i is finite. Under the interpretation as compact representations, then, there is no obstacle to conjoining the elements in each base layer H_i into a single sentence $h_i = \bigwedge H_i$. For the constructions we shall discuss in this paper, no change will result by such a maneuver. Rather

than \overrightarrow{H} , we can then equivalently use the string

$$\overrightarrow{h} = h_1 \prec \dots \prec h_n$$

It will be assumed throughout this paper that the beliefs of the highest priority, H_n or h_n , are consistent. Contradictions may arise only at lower levels.

Let \top (*verum*) and \perp (*falsum*) be the sentential constants that are always true and false, respectively. If we liked, we could extend strings by putting “ $\perp \prec$ ” in front or attaching “ $\prec \top$ ” to the end of \overrightarrow{h} . But the former is not necessary because, as we shall soon see, inconsistent “up-sets” are irrelevant anyway.¹ And the latter is not desirable as a general requirement on prioritized belief bases, because we want to allow revision methods that push up contingent sentences to the level of tautologies. As a consequence, the AGM postulate of ‘consistency preservation’, according to which only an inconsistent input can lead into an inconsistent belief set, is not validated by such methods studied in this paper.

We introduce some notation and abbreviations. Unless otherwise noted, i ranges from 1 to n :

$$\begin{aligned} H &= H_1 \cup \dots \cup H_n \\ H_{>i} &= H_i \cup \dots \cup H_n \\ H_{>i} &= H_{i+1} \cup \dots \cup H_n \text{ for } 0 \leq i \leq n-1 \\ h &:= h_1 \wedge \dots \wedge h_n \\ h_{\geq i} &:= h_i \wedge \dots \wedge h_n \\ h_{>i} &:= h_{i+1} \wedge \dots \wedge h_n \\ \overrightarrow{H_{>i}} &:= \langle H_i, \dots, H_n \rangle = H_i \prec \dots \prec H_n \\ \overrightarrow{H_{\leq i}} &:= \langle H_1, \dots, H_i \rangle = H_1 \prec \dots \prec H_i \\ \overrightarrow{h_{\geq i}} &:= h_i \prec \dots \prec h_n, \quad \overrightarrow{h_{\leq i}} := h_1 \prec \dots \prec h_i \\ \overrightarrow{h \wedge \alpha} &:= h_1 \wedge \alpha \prec \dots \prec h_n \wedge \alpha \\ \overrightarrow{h \vee \alpha} &:= h_1 \vee \alpha \prec \dots \prec h_n \vee \alpha \\ \overrightarrow{h \overset{+}{\vee} \alpha} &:= h_1 \prec h_1 \vee \alpha \prec h_2 \prec h_2 \vee \alpha \prec \dots \prec h_n \prec h_n \vee \alpha \end{aligned}$$

¹This statement has to be qualified. If the dynamics of belief are driven by syntactical manipulations on prioritized belief bases, it does matter whether there are lower levels that make the base inconsistent as a whole. Statically equivalent bases may be dynamically different. I neglect this point in the present paper.

And, for example

$$\begin{aligned} \overrightarrow{h_{\geq i} \wedge \alpha} &:= h_i \wedge \alpha \prec \cdots \prec h_n \wedge \alpha \\ \overrightarrow{h_{< i} \vee \alpha} &:= h_1 \vee \alpha \prec \cdots \prec h_{i-1} \vee \alpha \quad \text{for } 2 \leq i \leq n + 1 \\ \overrightarrow{h_{> i} \overset{+}{\vee} \alpha} &:= h_{i+1} \prec h_{i+1} \vee \alpha \prec \cdots \prec h_n \prec h_n \vee \alpha \quad \text{for } 0 \leq i \leq n - 1 \end{aligned}$$

For $\overrightarrow{h} = h_1 \prec \cdots \prec h_n$ and $\overrightarrow{g} = g_1 \prec \cdots \prec g_m$ we define the concatenations $\overrightarrow{h} \prec. \alpha = h_1 \prec \cdots \prec h_n \prec \alpha$ and $\overrightarrow{h} \prec. \overrightarrow{g} = h_1 \prec \cdots \prec h_n \prec g_1 \prec \cdots \prec g_m$. (The dot next to a \prec symbol indicates that at least one of the relata is not a set of sentences or a single sentence, but an ordered sequence itself.)

The most important sets definable by prioritized bases are the up-sets $H_{\geq i}$ and the sentences $h_{\geq i}$. They serve as standards of consistency and inconsistency in a way to be explained soon.

Instead of numbers, we can also use *sentences* in order to define the relevant up-sets. If H implies α , let in the following definitions i be the greatest number such that $H_{\geq i}$ implies α (so $H_{> i}$ does not imply α). Then we define $H_{\geq \alpha} = H_{\geq i}$ and $H_{> \alpha} = H_{> i}$ and $H_{= \alpha} = H_i$. If H does not imply α , we set $H_{\geq \alpha} = H_{> \alpha} = H$. Notice that $H_n = H_{=\top}$, but not necessarily $H_1 = H_{=h}$ (but this does hold for purified bases, see below). In the same fashion, we define $h_{\geq \alpha} = h_{\geq i}$, $h_{> \alpha} = h_{> i}$ and $h_{= \alpha} = h_i$, where i is the greatest number such that $h_{\geq i}$ implies α . Notational devices mixing sentences and numbers like $h_{> \alpha + 1}$ or $h_{= h - 1}$ should be understood in the obvious way.

The *belief set* \mathcal{B} supported by a prioritized base \overrightarrow{H} is defined as $Bel(\overrightarrow{H}) = Cn(H_{> \perp})$. Here and throughout this paper, we use Cn to indicate a consequence operation governing the language that is Tarskian, includes classical propositional logic and satisfies the deduction theorem.² Notice that belief sets so conceived are always consistent (except perhaps in the limiting case when H_n is itself inconsistent).

Beliefs in $Bel(\overrightarrow{H})$ can be ranked according to their certainty, reliability or importance. We employ a *Weakest Link Principle* according to which a chain is just as strong as its weakest link. Less metaphorically, a set of premises is just as strong as its weakest element. In accordance with this idea, it would be possible to define $rank_H(\alpha)$ to be the largest integer i

²A logic Cn is Tarskian iff it is reflexive ($H \subseteq Cn(H)$), monotonic (if $H \subseteq H'$, then $Cn(H) \subseteq Cn(H')$), idempotent ($Cn(Cn(H)) \subseteq Cn(H)$) and compact (if $\alpha \in Cn(H)$, then $\alpha \in Cn(H')$ for some finite $H' \subseteq H$). The deduction theorem says that $\alpha \rightarrow \beta \in Cn(H)$ if and only if $\beta \in Cn(H \cup \{\alpha\})$. We write $H \vdash \alpha$ for $\alpha \in Cn(H)$.

such that $H_{\geq i}$ implies α . But here is an important warning: *Numbers don't really mean anything in our framework* – never apply arithmetic operations (addition, subtraction, multiplication) to any such ranks! So let us work with a relation instead:

(Def \leq from \preceq) $\alpha \leq \beta$ iff for every i , if $H_{\geq i}$ implies α then it also implies β

Such relations are often called relations of *epistemic entrenchment*. The idea of (Def \leq from \preceq) has become folklore in the belief revision literature and was put to use, for instance by Rott (1991b) and Williams (1995). Entrenchment relations were first introduced and axiomatized by Gärdenfors and Makinson (1988). Notice, however, that the Gärdenfors-Makinson ‘maximality condition’ that says that only logical truths are maximally entrenched is not a necessary property of the entrenchment relations used in this paper.

An alternative and, as we said, more vivid representation of the significance of prioritized bases is in terms of possible worlds or more exactly, in terms of models of the underlying language. A prioritized base may be thought of as structuring the space of all models of the underlying language into a system $\$$ of nested spheres (à la Lewis 1973 and Grove 1988):

(Def $\$$ from \rightarrow): *The system of spheres $\$$ generated by a prioritized belief base \overrightarrow{H} is the set of sets S_i of models such that for each i , S_i is the set of models of $H_{\geq i}$, in symbols:*

$$\$ = \{ \text{mod}(H_{\geq i}) : i = 1, \dots, n \}$$

The idea is that the models of $H = H_{\geq 1}$ are the most plausible worlds, the models of $H_{\geq 2}$ that are not models of $H_{\geq 1}$ are the second most plausible worlds, etc., and the models of H_n that are not models of $H_{\geq n-1}$ are the least plausible worlds — except for those models that do not even satisfy H_n and may be regarded as completely ‘inaccessible’ to the agent’s mind. The set H_n characterizes the agent’s ‘certainties’ or ‘commitment set’ the elements of which he or she is extremely reluctant to give up.³

Now let us introduce an important operation on prioritized belief bases. Prioritized bases can be simplified or ‘purified’ without affecting \overrightarrow{H} the generated positions of the beliefs or worlds. The *purification* of a base \overrightarrow{H} deletes,

³Though not absolutely reluctant, see for instance the models of moderate and very radical expansion below. Segerberg (1998) unofficially calls what is characterized by H_n ‘knowledge’. It seems to me, however, that it is more adequate (though not fully adequate) to identify knowledge with belief that is indefeasible by true inputs, and this does not require maximal entrenchment. For discussions of this concept of knowledge in game-theoretic and epistemological contexts, see Stalnaker (1996) and Rott (2004), respectively.

for every i , all sentences α in H_i which are entailed by $H_{>i}$ (where $i < n$). If after these deletions a set H_i turns out to be empty, it is deleted as a coordinate from \overrightarrow{H} . Similarly, the purification of a base h deletes, for every i , every sentence h_i which is entailed by $h_{>i}$ (where again $i < n$). If some H_i or h_i is deleted, so is the symbol ' \prec ' to its right. Purification makes prioritized bases less misleading. If $h_i \prec h_j$ in a purified base, then it is guaranteed that also $h_i < h_j$ in the generated epistemic entrenchment ordering.⁴ If $\overrightarrow{H} = H_1 \prec \dots \prec H_n$ is purified, then the number of spheres in the generated system of spheres $\$$ is n , and the number of equivalence classes in the generated entrenchment relation \leq is $n + 1$ if \overrightarrow{H} is consistent, and n if \overrightarrow{H} is inconsistent. It is easy to see that the entrenchment relation or system of spheres generated by a purified base is identical with that of the unpurified one. For this reason we regard an original base and its purified forms as equivalent. We may always purify a prioritized base, but we are not forced to, when performing any of the belief change operations that follow.

The aim of this paper is to show that the most common qualitative approaches to iterated revision can be represented in a smooth, perspicuous and computationally efficient way as operations on prioritized bases. The operations have a much more constructive flavour than the equivalent operations on entrenchment relations or systems of models. In contrast to the latter, they are syntactic rather than semantic in nature.

Doxastic states \mathcal{S} can be represented, e.g., by systems of spheres of possible models $\$$, by entrenchment relations \leq or by prioritized belief bases \overrightarrow{H} . Doxastic states define belief sets, e.g., using the equations $K = Bel(\$) = \{\alpha : \alpha \text{ is true in all models that are contained in every non-empty } S \in \$\}$, $K = Bel(\leq) = \{\alpha : \perp < \alpha\}$ and $K = Bel(\overrightarrow{H}) = Cn(H_{>\perp})$, respectively.⁵ In the following, we indeed assume that K is derived from some state \mathcal{S} , i.e., from some system of spheres $\$$, some entrenchment relation \leq or some prioritized belief base \overrightarrow{H} . The traditional AGM notation $K * \alpha$ denoting the revised belief set is then to be read as an abbreviation for $Bel(\mathcal{S} * \alpha)$, i.e., $Bel(\$ * \alpha)$, $Bel(\leq * \alpha)$ or $Bel(\overrightarrow{H} * \alpha)$, respectively.

We now take revision operations to operate on doxastic states. The revisions below will be presented as revisions of \overrightarrow{h} rather than \overrightarrow{H} , but the

⁴Compare the 'Entailment condition' as a test for determining whether an 'E-base' is really an E-base for its generated entrenchment relation in Rott (1991b, p. 146).

⁵These definitions guarantee that the belief set is consistent, except in extreme limiting cases in which $\$$ contains only the empty set, $\top \leq \perp$, or H_n is inconsistent, respectively.

generalization to the latter case will be obvious. Just do ‘the same’ that is being done to h_i to all the members of H_i individually, and keep them together at their common level. Limitations of space not only make explicit proofs impossible, but also prevent me from dealing with the limiting cases in due detail. The reader is asked to read all coming claims and conditions for revisions by α as restricted to the case in which α is considered possible by the agent, i.e., in which h_n is consistent with α , or equivalently, in which some sphere in $\$$ contains some α -models, or again equivalently, in which $\neg\alpha$ is less entrenched than \top .

3. Variants of expansion

In traditional AGM theory, the expansion of a set of plain beliefs consists in simply adding a sentence to a given stock of beliefs and closing under deduction. This is a clear method offering no possibilities of choice. Its disadvantage becomes evident, however, when the input sentence is inconsistent with the prior beliefs. There is room for choice and different methods, however, if we consider expansions of belief states (like prioritized bases, entrenchment relations, system of spheres) rather than just belief sets. We think of expansions as applying sensibly to the paradigm case where the input is consistent with the prior beliefs, that is, with h . No claim is made that the expansion methods must make sense in the belief-contravening case. But it is instructive to study the number of possibilities of expansions (see Figs. 1–5) that mirror quite nicely the corresponding revision and contraction operations. No such analogy between (trivial) expansions and (non-trivial) revisions of belief sets is present in AGM theory.

Conservative expansion by α :	$\xrightarrow{\quad}$	h	\mapsto	$\alpha \prec. \xrightarrow{\quad} h$
Plain expansion by α :	$\xrightarrow{\quad}$	h	\mapsto	$h_1 \wedge \alpha \prec. \xrightarrow{\quad} h_{>1}$
Moderate expansion by α :	$\xrightarrow{\quad}$	h	\mapsto	$\xrightarrow{\quad} h \prec. \alpha \prec. \xrightarrow{\quad} h \vee \alpha$
Radical expansion by α :	$\xrightarrow{\quad}$	h	\mapsto	$\xrightarrow{\quad} h_{<n} \prec. h_n \wedge \alpha$
Very radical expansion by α :	$\xrightarrow{\quad}$	h	\mapsto	$\xrightarrow{\quad} h \prec. \alpha$

One can see immediately the symmetry between plain and radical and between conservative and very radical expansion. The differences are only those between inserting the input sentence *at* the lowest or highest level vs. inserting it in a *newly created* lowest or highest level of priority. It is also very evident now why the moderate method is called ‘moderate’: the input sentence gets assigned a middle rank. Let us have a look what happens to the number of different levels after purification of the revised prioritized

base. Assume the principal case for expansion in which h implies neither α nor $\neg\alpha$. Then plain expansion leaves the number of levels at n , while conservative expansion raises it to $n + 1$. Radical expansion give at most n levels, very radical expansion at most $n + 1$, and finally, moderate expansion gives at least $n + 1$ and at most $2n + 1$ levels. (Very) radical expansion tends to coarsen, while moderate expansion tends to refine it.

Very radical expansion accepts α as more certain than all the previous beliefs, thereby making some previously inaccessible α -models accessible. That is, some previously maximally entrenched sentences lose their status as certainties.⁶ This prevents very radical expansion from being commutative. In radical (but not very radical) expansion, the new information gets as highly entrenched as the maximal prior information, but not higher than that. This operation is commutative.

4. Radical revision

We take as paradigmatic for revision the case where the new information is incompatible with the original belief set (the *belief-contravening* case). We continue to assume that the agent is bound to accept the input sentence α and denote the posterior entrenchment relation by \leq' .

All methods for iterated revision to be discussed in this paper have essentially AGM revision as a limiting case for the case of a one-step revision. In terms of systems of spheres, this means that the innermost sphere of the revised SOS is exactly the intersection of the set of models of α with the smallest sphere in the original SOS that contains any models of α . In terms of entrenchments, a sentence β is more entrenched than \perp with respect to the revised entrenchment relation \leq' if and only if the conditional $\alpha \rightarrow \beta$ is more entrenched than $\neg\alpha$ with respect to the original entrenchment relation \leq .⁷

The SOS representations of each of the belief change operations to come are given in the Appendix. In the following main text, we list the corresponding operations on prioritized bases, the entrenchment representations, and finally the characterizations in terms of iterated belief changes.

The first method that we discuss is Segerberg's (1998) 'irrevocable revision' (see Fig. 6), which I like to call 'radical revision'. Fermé (2000) studied the same operation in terms of epistemic entrenchment.

⁶The same happens in 'moderate' expansion. Cf. footnote 10 below.

⁷In the limiting case in which no sphere in the SOS contains α -models, or in which $\neg\alpha$ is maximally entrenched, we can decide that the revised SOS and the revised entrenchment relation are identical with the original ones.

Here is a representation of the radical revision of \vec{h} by the input α .

$$\vec{h} \mapsto \vec{h}_{<n} \prec. h_n \wedge \alpha$$

The new base may then be purified. Equivalently, one could use the representation $h \wedge \alpha$ plus purification.

An even more radical strategy is the recipe of *very radical revision* (see Fig. 7):

$$\vec{h} \mapsto \vec{h} \prec. \alpha$$

Here the same comments apply as in the case of very radical expansion.

The revised entrenchment relation generated by the radical revision of a prioritized base is defined by

$$\gamma \leq' \delta \text{ iff } \alpha \rightarrow \gamma \leq \alpha \rightarrow \delta$$

the revised entrenchment relation generated by the very radical revision of a prioritized base is defined by

$$\gamma \leq' \delta \text{ iff } \alpha \rightarrow \gamma \leq \alpha \rightarrow \delta, \text{ and } \alpha \not\vdash \gamma \text{ or } \alpha \vdash \delta$$

The recipe for radical revision corresponds to the rule (RER) of Rott (1991a, p. 171; 2003, p. 130).⁸ This operation revises an arbitrary prior entrenchment ordering \leq without assuming that it was generated from a prioritized belief base.

Against the background of the AGM axioms for one-step revisions, radical revision can be characterized in terms of an iterated revision postulate as follows

$$(K * \alpha) * \beta = K * (\alpha \wedge \beta)$$

and very radical revision is similarly characterized by

$$(K * \alpha) * \beta = \begin{cases} K * (\alpha \wedge \beta) & \text{if } K * (\alpha \wedge \beta) \not\vdash \perp \\ Cn(\beta) & \text{if } \{\alpha, \beta\} \vdash \perp \\ Cn(\alpha, \beta) & \text{otherwise} \end{cases}$$

5. Conservative revision

Conservative revision (see Fig. 8), originally called ‘natural revision’, was advocated and studied by Boutilier (1993, 1996) and Rott (2003).

⁸Actually, (RER) in the later paper has an extra clause ‘and $\not\vdash \gamma$ or $\vdash \delta$ ’ that guarantees that $\alpha <' \top$ for non-tautological α . As already mentioned, I do not want to require AGM’s maximality condition in the present paper.

Here is a representation of the conservative revision of \overrightarrow{h} by the input α .

$$\overrightarrow{h} \mapsto \alpha \prec . \overrightarrow{h_{\leq \neg\alpha}} \vee \alpha \prec . \overrightarrow{h_{> \neg\alpha}}$$

If h is purified and does not imply α , no posterior purification is necessary, and the posterior base has $n + 1$ levels. If h is purified and does imply α , then the term ‘ $\alpha \prec$ ’ will be dropped in purification, and the posterior base is identical with the n -level prior base \overrightarrow{h} .

The revised entrenchment relation \leq' generated by the conservative revision of a prioritized base by α is defined by

$$\gamma \leq' \delta \text{ iff } \alpha \rightarrow \gamma \leq \neg\alpha, \text{ or } \gamma \leq \delta \text{ and } \neg\alpha < \alpha \rightarrow \delta$$

This is the condition (CER) for ‘conservative entrenchment revision’ of Rott (2003, p. 122).

Conservative revision can be characterized in terms of an iterated revision postulate as follows

$$(K * \alpha) * \beta = \begin{cases} K * (\alpha \wedge \beta) & \text{if } \beta \text{ is consistent with } K * \alpha^9 \\ K * \beta & \text{otherwise} \end{cases}$$

Mirroring the difference between conservative and plain (AGM) expansion, we can define a variant of conservative revision which is obtained by an AGM contraction (see Section 8 below) with respect to $\neg\alpha$, followed by a plain (AGM) expansion by α , i.e., by a version of the well-known Levi identity:

Here is a representation of what might be called the *plain revision* of \overrightarrow{h} by the input α .

$$\overrightarrow{h} \mapsto \alpha \prec . \overrightarrow{h_{>1, \leq \neg\alpha}} \vee \alpha \prec . \overrightarrow{h_{> \neg\alpha}}$$

This operation forgets as it were about the lowest ranked elements of the prioritized belief base, or correspondingly, the innermost ring of the prior system of spheres. Since it does not seem to be a very natural revision operation, I refrain from giving alternative representations of it.

6. Moderate revision

Moderate revision is my name for what is often called ‘lexicographic revision’ (see Fig. 9). It has been advocated and studied by Nayak and his collabo-

⁹In a more general context without ‘dispositional coherence’, we should put $Cn((K * \alpha) \cup \{\beta\})$ in this case rather than $K * (\alpha \wedge \beta)$, see Rott (2003). But given the dispositional coherence encoded in AGM’s 7th and 8th axioms, this comes down to the same thing.

rators (1994, 2003), but also by many other researchers. It has become part of the folklore of belief revision research, but here does not seem to be a standard reference paper for it. We present a formulation here that does not presume consistency preservation for revision functions (or the maximality condition for entrenchments).

Here is a representation of the moderate revision of \vec{h} by the input α .

$$\vec{h} \mapsto \vec{h} \prec . \alpha \prec . \overrightarrow{h \vee \alpha}$$

As always, the new base may be purified.

The revised entrenchment relation \leq' generated by the moderate revision of a prioritized base by α is defined by

$$\gamma \leq' \delta \text{ iff } \begin{cases} \alpha \rightarrow \gamma \leq \alpha \rightarrow \delta \text{ and } \alpha \not\vdash \gamma & \text{or} \\ \gamma \leq \delta \text{ and } \alpha \vdash \delta \end{cases}$$

This is similar to, but not exactly the same as condition (MER) (for ‘moderate entrenchment revision’) of Rott (2003, p. 131). The slight modification suggested here is correct also when the revision function does not satisfy the fifth AGM postulate (‘consistency preservation’) or, equivalently, when the entrenchment relation has not only tautologies as maximal elements (what was excluded in Rott 2003).

Moderate revision can be characterized in terms of an iterated revision postulate as follows¹⁰

$$(K * \alpha) * \beta = \begin{cases} K * (\alpha \wedge \beta) & \text{if } K * (\alpha \wedge \beta) \text{ is consistent} \\ K * \beta & \text{if } \alpha \vdash \neg\beta \\ Cn(\alpha \wedge \beta) & \text{otherwise} \end{cases}$$

7. Restrained revision

Recently, Booth and Meyer (2006) advocated the interesting operation of restrained revision (see Fig. 10), which can be seen as composition of a refinement by α (see Section 9) followed by a conservative revision by α .

¹⁰A somewhat *more moderate revision* could be defined thus:

$$\vec{h} \mapsto \vec{h} \prec . \alpha \prec . \overrightarrow{h_{<n} \vee \alpha} \prec . h_n$$

In terms of SOS’s, the more moderate revision never turns previously inaccessible worlds into accessible ones (what moderate revision usually does). In terms of entrenchments, more moderate revision replaces the clauses ‘ $\alpha \not\vdash \gamma$ ’ and ‘ $\alpha \vdash \delta$ ’ by ‘ $\alpha \rightarrow \gamma < \top$ ’ and ‘ $\top \leq \alpha \rightarrow \delta$ ’, respectively. In terms of iterated revision, more moderate revision replaces the last two lines of moderate revision by ‘ $(K*\alpha)*\beta = K*\beta$ otherwise’. Also cf. footnote 15.

Here is a representation of the restrained revision of \vec{h} by the input α .

$$\vec{h} \mapsto \alpha \prec. \overrightarrow{h_{\leq -\alpha} \vee \alpha} \prec. \overrightarrow{h_{> -\alpha} \dot{\vee} \alpha}$$

plus purification.

The revised entrenchment relation \leq' generated by the restrained revision of a prioritized base by α is defined by

$$\gamma \leq' \delta \text{ iff } \alpha \rightarrow \gamma \leq \neg\alpha, \text{ or } \gamma \leq \delta \text{ and } \begin{cases} \alpha \rightarrow \gamma \leq \gamma \\ \neg\alpha < \alpha \rightarrow \delta \text{ and } \gamma < \alpha \rightarrow \delta \end{cases} \text{ or}$$

While its entrenchment representation is somewhat difficult to comprehend, restrained revision can be characterized elegantly in terms of an iterated revision postulate (Booth and Meyer 2006):

$$(K * \alpha) * \beta = \begin{cases} K * (\alpha \wedge \beta) & \text{if } K * \alpha \not\vdash \neg\beta \text{ or } K * \beta \not\vdash \neg\alpha \\ K * \beta & \text{otherwise} \end{cases}$$

8. Variants of contraction

The simplest way of getting rid of a belief α is a method that has been called ‘Rott contraction’ by Fermé and Rodriguez (1998), ‘severe withdrawal’ by Pagnucco and Rott (1999) and ‘mild contraction’ by Levi (2004) (see Fig. 11). The method was extended to iterated belief change in Rott (2006).

Here is a representation of the *severe withdrawal* of α from \vec{h}

$$\vec{h} \mapsto \overrightarrow{h_{> \alpha}}$$

The revised entrenchment relation corresponding to the severe withdrawal operation with respect to α is this:

$$\gamma \leq' \delta \text{ iff } \gamma \leq \alpha \text{ or } \gamma \leq \delta$$

The so-called Levi identity recommends to construct a revision by α through applying an operation of expansion by α after a preparatory contraction by $\neg\alpha$. Accordingly, we can define different concepts of *severe revision* by applying different expansion operations after a severe withdrawal. Let us distinguish three versions of severe revision by α .

This is severe withdrawal combined with conservative expansion (see Fig. 12):

$$\vec{h} \mapsto \alpha \prec. \overrightarrow{h_{> \neg\alpha}}$$

Here is severe withdrawal combined with plain expansion (see Fig. 13):

$$\overrightarrow{h} \mapsto \overrightarrow{h_{=-\alpha+1} \wedge \alpha} \prec. \overrightarrow{h_{>-\alpha+1}}$$

And here is severe withdrawal combined with moderate expansion (see Fig. 14):

$$\overrightarrow{h} \mapsto \overrightarrow{h_{>-\alpha}} \prec. \alpha \prec. \overrightarrow{h_{>-\alpha} \vee \alpha}$$

The most faithful extrapolation of one-step *AGM contraction* of belief sets to the revision of belief states seems to be the *conservative contraction* with respect to α (see Fig. 15):

$$\overrightarrow{h} \mapsto \overrightarrow{h_{\leq \alpha} \vee \neg \alpha} \prec. \overrightarrow{h_{> \alpha}}$$

The revised entrenchment relation corresponding to the conservative contraction operation with respect to α is this:

$$\gamma \leq' \delta \quad \text{iff} \quad \gamma \leq \perp \quad \text{or} \quad \alpha \vee \gamma \leq \alpha \quad \text{or} \quad (\alpha < \alpha \vee \delta \quad \text{and} \quad \gamma \leq \delta)$$

Notice that $\gamma \leq' \delta$ according to conservative contraction implies $\gamma \leq' \delta$ according to severe withdrawal.

Finally, here is a representation of the *moderate contraction* (see Fig. 16) of \overrightarrow{h} with respect to α .

$$\overrightarrow{h} \mapsto \overrightarrow{h_{> \alpha}} \prec. \overrightarrow{h \vee \neg \alpha}$$

Nayak, Goebel and Orgun (2007) propose an operation of *lexicographic contraction* which corresponds to Nayak and others' operation of lexicographic revision. This interesting proposal is, however, too complex to receive a treatment in the present paper.

9. Refinement: Neither revision nor contraction

Papini (2001) introduced an interesting belief change operation that is neither a revision nor a contraction operation. It is a kind of *reverse lexicographic belief change*, that I like to call *refinement*. In the system of spheres modelling, each level of the prior system is kept in place, but split in such a way that the α -models of a certain level are after the change made more plausible than the $\neg\alpha$ -models of the same level (see Fig. 17).

Refinement of \overrightarrow{h} by input α .

$$\overrightarrow{h} \mapsto \overrightarrow{h_{<-\alpha}} \prec. \overrightarrow{h_{\geq-\alpha} \overset{+}{\vee} \alpha}$$

plus purification.

The revised entrenchment relation \leq' generated by the reverse lexicographic change of a prioritized base by α is defined by

$$\gamma \leq' \delta \text{ iff } \gamma \leq \delta \text{ and } \begin{cases} \alpha \rightarrow \gamma \leq \gamma & \text{or} \\ \gamma < \alpha \rightarrow \delta \end{cases}$$

Using the notation K/α for the belief set resulting from the refinement of \overrightarrow{H} or \leq or $\$$ by α , we note that the operation $/$ is not always successful in the way revision operations $*$ are supposed to be successful. More precisely, we have $\alpha \in K/\alpha = Cn(K \cup \{\alpha\})$ if and only if α is consistent with K ; otherwise $\neg\alpha \in K/\alpha = K$.

There is no characterization of reverse lexicographic belief change ('refinement') in terms of iterated 'revision' postulates, perhaps simply because refinement is no revision operation.¹¹ Refinement need not have any effects on the belief set level, but may be confined to worlds in outer systems of spheres or to sentences higher up in the entrenchment ranking. We have the property (which is too weak to characterize refinement)

$$K/\alpha/\beta = \begin{cases} K/(\alpha \wedge \beta) = K + (\alpha \wedge \beta) & \text{if } \neg(\beta \wedge \alpha) \notin K \\ K/\alpha = K + \alpha & \text{if } \neg\alpha \notin K, \neg(\beta \wedge \alpha) \in K \\ K/\beta = K + \beta & \text{if } \neg\alpha \in K, \neg\beta \notin K \\ K/\alpha = K/\beta = K/(\alpha \wedge \beta) = K & \text{if } \neg\alpha, \neg\beta \in K \end{cases}$$

10. Two-dimensional operators: Revision by comparison

The idea of two-dimensional belief change operators is that a belief state is transformed in such a way that a sentence α (the 'input') gets accepted with the certainty of a sentence β (the 'reference sentence'). The input is something like ' $\beta \leq \alpha$ '. The operation of revision by comparison (see Fig. 18) was studied by Cantwell (1997), who called it 'raising', and by Fermé and Rott (2004), who used the notation $\circ_\beta \alpha$. The principal case is when β is more entrenched than α (which we may think of not being accepted in the prior belief state); some interesting limiting cases will be addressed presently.

$$\overrightarrow{h} \quad \mapsto \quad \overrightarrow{h}_{<\beta} \prec. h_{=\beta} \wedge \alpha \prec. \overrightarrow{h}_{>\beta}$$

plus purification.

Fermé and Rott (2004, p. 13) give the following definition of revision by comparison in terms of epistemic entrenchment. Let \leq be a prior entrenchment ordering (usually not thought of arising from an e-base). Assuming

¹¹A characterization should be possible using a postulate for $K/\alpha * \beta$, where $*$ is an AGM revision function.

again that the agent is to accept α (the input sentence) at least as certainly as β (the reference sentence), the posterior entrenchment relation $\leq' = \leq_{\beta \leq \alpha}^*$ is defined by

$$\gamma \leq' \delta \text{ iff } \begin{cases} \beta \wedge (\alpha \rightarrow \gamma) \leq (\alpha \rightarrow \delta) \text{ and } \gamma \leq \beta & \text{or} \\ \gamma \leq \delta \text{ and } \beta < \gamma \end{cases}$$

It is surprising that the extremely simple operation on prioritized bases indeed captures the operation of revision by comparison which was characterized and studied in rather laborious ways by Fermé and Rott.

There are a number of interesting *unary special cases of revision by comparison*. The special case $\circ_{\alpha} \perp$ with input sentence \perp and reference sentence α reduces to a severe withdrawal of α (cf. Section 8). The special case $\circ_{\top} \alpha$ with input sentence α and reference sentence \top reduces to an irrevocable or radical revision by α (cf. Section 4). Another operation worth mentioning is that of irrefutable revision obtained by fixing a reference sentence ε and defining $K * \alpha := K \circ_{\varepsilon} \alpha$. Though similar with irrevocable revision, especially if a highly entrenched reference sentence ε is chosen, there are some interesting differences (cf. Rott 2006). The change of the prioritized knowledge base cannot further be reduced (see Fig. 19).

11. Two-dimensional operators: Cantwell’s lowering

Cantwell (1997) argued that there are two ways of dealing with the situation when we have the prior relation $\alpha < \beta$ and when the input is something like ‘ $\beta \leq \alpha$ ’. What ‘revision by comparison’ in the sense of Fermé and Rott does in some intuitive way is to promote α to the rank of β . Although it is problematic to make cross-relational comparisons,¹² the above representation with prioritized belief bases illustrates this: α is simply inserted into the rank of β . But Cantwell saw that there is also a dual operation. One can also obtain the intended effect, in the same principal situation, by demoting β to the rank of α (see Fig. 20).¹³ This is the relevant operation on prioritized bases:

$$\xrightarrow{h} \quad \mapsto \quad \xrightarrow{h_{<\alpha}} \prec. h_{=\alpha, \leq \beta} \wedge \beta \prec. \xrightarrow{h_{>\alpha, \leq \beta} \vee \neg \beta} \prec. \xrightarrow{h_{>\beta}}$$

plus purification.

¹²Compare Fermé and Rott (2004, pp. 25–26).

¹³Something like that can happen in ‘revision by comparison’ as well, see the case of severe withdrawal above. However, the paradigm case for the application of revision by comparison is $\alpha \leq \perp < \beta$, while the paradigm case of lowering is $\perp < \alpha < \beta$.

Assuming that the agent is to accept α (the input sentence) at least as certainly as β (the reference sentence) and that $\beta < \top$, the revised entrenchment relation $\leq' = \leq_{\beta < \alpha}^*$ as generated by the lowering of β to the degree of α is defined by the following recipe:¹⁴

$$\gamma \leq' \delta \text{ iff } (\gamma \leq \delta \text{ and } \gamma \leq \alpha) \text{ or } (\gamma \leq \delta \text{ and } \beta < \beta \vee \delta) \text{ or } (\alpha \leq \delta \text{ and } \beta \vee \gamma \leq \beta)$$

As far as I know, this condition is new, but it is similar in spirit to Cantwell’s axiomatization of lowering. It looks more complicated than it is. Roughly, the explanation for it is this: The old ordering \leq remains undisturbed below α , and indeed the relationship $\gamma \leq \delta$ does not change as long as δ is not lowered (which happens when $\beta < \beta \vee \delta$). If not $\gamma \leq \delta$, we can get a new relationship $\gamma \leq' \delta$ if γ is lowered (which happens when $\beta \vee \gamma \leq \beta$) and δ is at least as entrenched as α .

The above condition does not give us the lowering operation if $\top \leq \beta$, for in this case it reduces to $\gamma \leq' \delta$ iff $\alpha \leq \delta$ or $\gamma \leq \delta$. This operation is a kind of dual to severe withdrawal (which rules that $\gamma \leq' \delta$ iff $\gamma \leq \alpha$ or $\gamma \leq \delta$). While, roughly speaking, severe withdrawal collapses the levels below α into one, this operation collapses the levels above α into one, and in fact into the highest possible one (see Fig. 21).

Given our general assumption that $\alpha < \beta$ in the prior belief state, there are two *unary special cases of the lowering operation*. First, fix $\beta = \top$. Then the lowering operation on bases gives

$$\overrightarrow{h} \mapsto \overrightarrow{h_{<\alpha}} \prec. h_{=\alpha} \wedge \top \prec. \overrightarrow{h_{>\alpha} \vee \neg \top}$$

which produces no change at all. Tautologies simply cannot be lowered.

Second, fix $\alpha = \perp$. In this case the lowering operation reduces to

$$\overrightarrow{h} \mapsto h_{=\perp} \wedge \beta \prec. \overrightarrow{h_{>\perp, \leq \beta} \vee \neg \beta} \prec. \overrightarrow{h_{>\beta}}$$

which results in a conservative contraction (= AGM contraction) with respect to β .

The recipe for lowering is quite similar to the recipe for conservative revision. The obvious input that might show that conservative revision by α is in fact a special case of lowering, namely an extreme lowering of $\neg\alpha$, would

¹⁴In the case $\top \leq \beta$, the revised entrenchment relation $\leq' = \leq_{\beta < \alpha}^*$ generated by lowering is defined as

$$\gamma \leq' \delta \text{ iff } (\gamma \leq \delta \text{ and } \gamma \leq \alpha) \text{ or } (\gamma \leq \delta \text{ and } \vdash \beta \vee \delta) \text{ or } (\alpha \leq \delta \text{ and } \nVdash \beta \vee \gamma)$$

If even $\vdash \beta$, this reduces to $\gamma \leq' \delta$ iff $\gamma \leq \delta$.

be ‘ $\neg\alpha \leq \perp$ ’. But this doesn’t quite give us a revision, it only amounts to conservative contraction with respect to $\neg\alpha$. It eliminates $\neg\alpha$, but it does not promote α to the rank of a belief, i.e., above \perp . For revision, we need another kind of lowering operation. Inputs in the form of strict inequalities will help us to solve the problem (see Section 13 below).

12. Gentle raising and lowering

Revision by comparison (raising) is different from lowering even when α and β are ‘neighbours’ in the sense that in the prior entrenchment ordering, there is no sentence strictly between α and β . One can see this in the operations of ‘gentle promotion’ and ‘gentle demotion’, in which the rank of α is raised and lowered by one, respectively (see Figures 22 and 23).

Gentle promotion of α :

$$\vec{h} \mapsto \vec{h}_{\leq\alpha} \prec. h_{=\alpha+1} \wedge \alpha \prec. \vec{h}_{>\alpha+1}$$

Gentle demotion of α :

$$\vec{h} \mapsto \vec{h}_{<\alpha-1} \prec. h_{=\alpha-1} \wedge \alpha \prec. h_{=\alpha} \vee \neg\alpha \prec. \vec{h}_{>\alpha}$$

The reader is invited to compare this with the related operation advocated by Darwiche and Pearl (1997, p. 15).

13. Two-dimensional operators: Raising and lowering by strict comparisons

Now suppose the initial situation is that $\alpha \leq \beta$. Can an inequality $\beta < \alpha$ as input be processed in just the same way as the equality $\beta \leq \alpha$? First we have to be clear that not any input of the form $\beta < \alpha$ is admissible. If α implies β , then α cannot be more entrenched than β . So let us assume that α does not imply β , i.e., that $\alpha \rightarrow \beta$ is not a logical truth, and have a look at the official definitions.

Raising with input $\beta < \alpha$ would seem to be simply this:

$$\vec{h} \mapsto \vec{h}_{\leq\beta} \prec. h_{=\beta+1} \wedge \alpha \prec. \vec{h}_{>\beta+1}$$

But there is a precondition here if the operation is to be successful: $h_{>\beta} \wedge \alpha$ must not imply β . If it does, one has to put α somewhere further up, and exactly to the lowest level i such that $h_{>i} \wedge \alpha$ does not imply β . So the right idea to meet the constraint specified by the input is this (see Fig. 24):

$$\vec{h} \mapsto \vec{h}_{\leq(\alpha\rightarrow\beta)} \prec. h_{=(\alpha\rightarrow\beta)+1} \wedge \alpha \prec. \vec{h}_{>(\alpha\rightarrow\beta)+1}$$

It is clear that the new prioritized base generates the relation $\beta < \alpha$. The two sentences occupy neighbouring layers of entrenchment separated at the left occurrence of ‘ $\prec.$ ’.

Prima facie, *lowering* with input $\beta < \alpha$ would seem to be this:

$$\overrightarrow{h} \mapsto \overrightarrow{h_{<\alpha-1}} \prec. h_{=\alpha-1} \wedge \beta \prec. \overrightarrow{h_{\geq\alpha, \leq\beta} \vee \neg\beta} \prec. \overrightarrow{h_{>\beta}}$$

Due to the disjuncts ‘ $\neg\beta$ ’, there is no danger that β is implied by higher levels. However, here we have a problem complementary to the one before. It is no longer guaranteed that the levels higher than β after the change still imply α . The solution is similar. Again we have to replace β by $\alpha \rightarrow \beta$. The right recipe turns out to be, after a little simplification (see Fig. 25):

$$\overrightarrow{h} \mapsto \overrightarrow{h_{<\alpha-1}} \prec. h_{=\alpha-1} \wedge \beta \prec. \overrightarrow{h_{\geq\alpha, \leq(\alpha\rightarrow\beta)} \vee (\alpha \wedge \neg\beta)} \prec. \overrightarrow{h_{>(\alpha\rightarrow\beta)}}$$

It is clear that the new prioritized base generates the relation $\beta < \alpha$. Again the two sentences occupy neighbouring layers of entrenchment separated at the middle occurrence of ‘ $\prec.$ ’.

Conservative revision is a special case of lowering with strict inputs. The input is simply $\perp < \alpha$. It turns out that shifting contradictions below the level of α is nothing but conservatively accepting α .

14. Two-dimensional operators: Bounded revision

Conservative revision was soon recognized as being too conservative: Only *very few* α -models are made more plausible than the $\neg\alpha$ -models. On the other hand, moderate revision is still fairly radical: *All* α -models are treated as more plausible than *all* the $\neg\alpha$ -models. It seems a good idea to employ a two-dimensional operator to steer a middle course. Revision by comparison (raising) and lowering, however, are not the right solutions to this problem, since they are not “between” the one-dimensional operators of conservative and moderate revision, and they do not satisfy the Darwiche-Pearl postulates. *Bounded revision* is a two-dimensional operator that is in a precise sense between conservative and moderate revision. It is motivated and explored in Rott (2007). It seems to dispel Spohn’s (1988, pp. 112–113) early complaints about the disadvantages of both conservative and moderate revision.

14.1. Bounded revision, strict version.

The idea of this operation is to accept an input sentence α as long as β holds along with α (see Fig. 26). The reference sentence β functions here as a measure of how much of α the agent should consider most plausible after the change. This idea is similar to that of revision by comparison, but not quite the same. The operation to be considered in this subsection does not

move α to an entrenchment level exceeding that of β . (It is only the variant that we are going to consider in subsection 14.2 that achieves this.)

Here is a representation of the strict bounded revision of \overrightarrow{h} by input α as long as β (along with α).

$$\overrightarrow{h} \mapsto \overrightarrow{h} \prec . \alpha \prec . \overrightarrow{h_{<(\alpha \rightarrow \beta)}} \vee \alpha \prec . \overrightarrow{h_{\geq(\alpha \rightarrow \beta)}}$$

or equivalently, modulo purification,

$$\overrightarrow{h} \mapsto \overrightarrow{h_{>-\alpha, <(\alpha \rightarrow \beta)}} \prec . \alpha \prec . \overrightarrow{h_{<(\alpha \rightarrow \beta)}} \vee \alpha \prec . \overrightarrow{h_{\geq(\alpha \rightarrow \beta)}}$$

Now let us look at the definition of bounded revision in terms of epistemic entrenchment. If \leq is the prior entrenchment ordering, then the posterior entrenchment relation $\leq' = \leq_{\alpha; \beta}^*$ is given by

$$\gamma \leq' \delta \text{ iff } \begin{cases} \alpha \rightarrow \gamma \leq \alpha \rightarrow \delta & , \text{ if } \alpha \rightarrow (\gamma \wedge \delta) < \alpha \rightarrow \beta \\ \gamma \leq \delta & , \text{ otherwise} \end{cases}$$

In the following equation for iterated revisions, read $K * \alpha := K *_{;\varepsilon} \alpha$ and $K * \beta := K *_{;\varepsilon'} \beta$ for some ε and ε' .

$$(K * \alpha) * \beta = \begin{cases} K * (\alpha \wedge \beta) & \text{if } \neg(\alpha \wedge \beta) < \alpha \rightarrow \varepsilon \\ K * \beta & \text{otherwise} \end{cases}$$

Notation: Here $\neg(\alpha \wedge \beta) < \alpha \rightarrow \varepsilon$ is short for the condition that ε is in, but $\neg\beta$ is not in $K * (\alpha \wedge (\neg\beta \vee \varepsilon))$. This abbreviation is in accordance with usual entrenchment theories.

The strict version of bounded revision reduces to *moderate revision* if one takes a logical truth like \top as the reference sentence, except for a limiting case.¹⁵

14.2. Bounded revision, non-strict version.

The idea of this operation is to accept an input sentence α as long as β holds along with α , and even just a little more (see Fig. 27). The operation of this subsection moves α to an entrenchment level just above that of β . In this respect it is quite close to revision by comparison.

Here is a representation of the non-strict bounded revision of \overrightarrow{h} by input α as long as β (along with α).

$$\overrightarrow{h} \mapsto \overrightarrow{h} \prec . \alpha \prec . \overrightarrow{h_{\leq(\alpha \rightarrow \beta)}} \vee \alpha \prec . \overrightarrow{h_{>(\alpha \rightarrow \beta)}}$$

¹⁵The difference in the limiting case is precisely that between moderate and more moderate revision. See footnote 10.

or equivalently, modulo purification,

$$\overrightarrow{h} \mapsto \overrightarrow{h_{\geq \neg\alpha, \leq (\alpha \rightarrow \beta)}} \prec \cdot \alpha \prec \cdot \overrightarrow{h_{\leq (\alpha \rightarrow \beta)} \vee \alpha} \prec \cdot \overrightarrow{h_{> (\alpha \rightarrow \beta)}}$$

We again look at the definition of this version of bounded revision in terms of epistemic entrenchment. Let \leq be a prior entrenchment ordering. Then the posterior entrenchment relation $\leq' = \leq_{\alpha, \beta}^*$ is given by

$$\gamma \leq' \delta \text{ iff } \begin{cases} \alpha \rightarrow \gamma \leq \alpha \rightarrow \delta, & \text{if } \alpha \rightarrow (\gamma \wedge \delta) \leq (\alpha \rightarrow \beta) \\ \gamma \leq \delta, & \text{otherwise.} \end{cases}$$

In the following equation, read $K * \alpha := K *_{\varepsilon} \alpha$ and $K * \beta := K *_{\varepsilon'} \beta$ for some ε and ε' .

Iterated revision postulate

$$(K * \alpha) * \beta = \begin{cases} K * (\alpha \wedge \beta), & \text{if } \neg(\alpha \wedge \beta) \leq \alpha \rightarrow \varepsilon \\ K * \beta, & \text{otherwise.} \end{cases}$$

Notation: Here $\neg(\alpha \wedge \beta) \leq \alpha \rightarrow \varepsilon$ is short for the condition that either $\neg\beta$ is not in or ε is in $K * (\alpha \wedge (\neg\beta \vee \varepsilon))$. This abbreviation is in accordance with usual entrenchment theories.

The non-strict version of bounded revision reduces to *conservative revision* if one takes $\neg\alpha$ as the reference sentence, except for a limiting case.¹⁶

15. Conclusion

A prioritized belief base represents an agent’s belief state. The set of her beliefs as well as her ranking of beliefs in terms of entrenchment can easily be obtained from a prioritized base. The prioritized base representation has, I believe, a number of significant advantages over the more established models. It is compact, constructive and convenient. While the semantics of spheres of possible worlds helps us understand the changes of belief states very well, the syntax of prioritized bases helps us to read off at a glance much of the contents and ranks of a base. We have used bases as compact and convenient tools for representing belief states, without implying that the elements of such a base themselves carry any epistemological weight as “basic” or “explicit” beliefs. Bases are finite and typically have only a comparatively small number of layers and a small number of sentences within

¹⁶In terms of iterated revision, for instance, the difference is as follows. Non-strictly bounded revision with $\varepsilon = \neg\alpha$ gives the inconsistent set $K * \alpha * \beta = K * (\alpha \wedge \beta)$ if $K * \alpha$ is inconsistent, while our official definition of conservative revision gives $K * \alpha * \beta = K * \beta$ in this case. This difference could easily be adapted, if we liked.

each layer. In contrast, other representations of doxastic states typically involve large numbers of possible worlds, or of beliefs to be ordered by some preference relation.

We have presented a fairly wide, though certainly not exhaustive, variety of methods for belief revision by way of manipulations of prioritized bases. These manipulations display quite clearly where in an existing priority ordering the new input is being placed: at the bottom (conservative revision, severe revision), at the top (radical revision) or somewhere in the middle (moderate revision, raising and lowering). There is a surprising multiplicity of revision methods that can be captured in this way. We have collected sphere models of 27 change functions in the Appendix.

A main point of this paper has been to show that prioritized bases are a very good way of representing not only belief states at a certain time, but also changes of belief states. Besides the calculation of implications, the operations to be performed on prioritized bases are: Copying some list of sentences, cutting some such list, applying booleans (\neg , \vee and \wedge) to the elements of a list, and concatenating lists. Prioritized belief base engineering is a little like DNA engineering. It probably is not realistic psychologically, but it should have nice computational properties. All operations are simple, transparent and give the user an immediate feeling of the status that a new piece of input is assigned in the posterior belief state.

We have gathered considerable inductive evidence that the revisions of belief states systematically defined via SOSs can all be captured by fairly simple syntactical means (prioritized belief base engineering). It does not seem that this can be proved, however, given the vagueness of the terms “systematically defined” and “fairly simple”.

After this paper was conceived, I was alerted to the fact that ideas very similar to prioritized base changes as presented here have already been explored in the framework of possibilistic logic in a series of papers, e.g. in Benferhat, Dubois and Prade (2001). The work of Meyer, Ghose and Chopra (2001) is also relevant. The research of both groups was done in the more general (and more interesting) area of belief merging. I recommend the reader to closely study these works and also consult the references mentioned therein. The present paper complements these earlier works in the following respects. The presentation as given here is somewhat simpler; I survey a larger number of methods of iterated belief change that can all lay claim to being regarded as rational; and finally, I make it fully clear that no numbers are needed for any of the belief change methods considered.

All the methods considered are purely qualitative, in the sense that there are no meaningful numbers involved. The numbers used in the representa-

tion of prioritized belief bases, as well as the numbers appearing in the sphere pictures only encode orderings. In view of the abundance of qualitative methods at our disposal, we are not likely to subscribe to the view of proponents of numerical methods, according to which purely qualitative methods will always remain too poor to model a reasonable evolution of our beliefs. The problem is rather the reverse: We are facing an embarrassment of riches. What we urgently need is some substantive metatheory that tells us which method to apply in what situations. Unfortunately, we do not have anything like such a methodology yet.

Acknowledgements. This paper is a corrected and improved version of a paper with the same title that appeared in Henrik Lagerlund, Sten Lindström and Rysiek Sliwinski (eds.), *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, Uppsala Philosophical Studies, Vol. 53, University of Uppsala 2006. First versions of this paper were presented at Schloss Dagstuhl and the Certosa di Pontignano, Siena, in 2005. I am grateful to the organizers Jim Delgrande, Jérôme Lang, Cristiano Castelfranchi and Fabio Paglieri for the opportunity to give the paper there, and to the audiences for excellent discussions. I thank Didier Dubois and Henri Prade for drawing my attention to the existence of related work in possibilistic logic, Eduardo Fermé, Brian Hill, Franz Huber, and two anonymous referees for helpful comments on earlier versions of this paper, and Ralf Busse for making me aware that Peirce's beautiful quote on diagrammatic reasoning is relevant for the theory of belief revision. Last but not least I want to thank David Makinson, Jacek Malinowski and Heinrich Wansing for encouraging me to submit this paper for inclusion in the present volume.

References

- ALCHOURRÓN, CARLOS, GÄRDENFORS, PETER, MAKINSON, DAVID (1985), 'On the logic of theory change: Partial meet contraction functions and their associated revision functions', *Journal of Symbolic Logic*, 50: 510–530,
- ARECES, CARLOS, BECHER, VERONICA (2001), 'Iterable AGM functions', in Williams, Rott (eds.), pp. 261–277.
- BENFERHAT, SALEM, DUBOIS, DIDIER, PRADE, HENRI (2001), 'A computational model for belief change and fusing ordered belief bases', in Williams, Rott (eds.), pp. 109–134.
- BOOTH, RICHARD, MEYER, THOMAS (2006), 'Admissible and restrained revision', *Journal of Artificial Intelligence Research*, 26: 127–151.
- BOUTILIER, CRAIG, (1993), 'Revision sequences and nested conditionals', in Bajcsy, R. (ed.), *IJCAI-93—Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pp. 519–525.

- BOUTILIER, CRAIG (1996), 'Iterated revision and minimal change of conditional beliefs', *Journal of Philosophical Logic*, 25: 263–305.
- CANTWELL, JOHN (1997), 'On the logic of small changes in hypertheories', *Theoria*, 63: 54–89.
- DARWICHE, ADNAN, PEARL, JUDEA (1997), 'On the logic of iterated belief revision', *Artificial intelligence*, 89: 1–29.
- DUBOIS, DIDIER, LANG, JÉRÔME, PRADE, HENRI, 'Possibilistic logic', in: Gabbay, D.M., Hogger, C.J., Robinson, J.A. (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3, Nonmonotonic Reasoning and Uncertain Reasoning*, Clarendon Press, Oxford, 1994, pp. 439–513.
- FERMÉ, EDUARDO (2000), 'Irrevocable belief revision and epistemic entrenchment', *Logic Journal of the IGPL*, 8: 645–652.
- FERMÉ, EDUARDO, RODRIGUEZ, RICARDO (1998), 'A brief note about Rott contraction', *Logic Journal of the IGPL*, 6: 835–842.
- FERMÉ, EDUARDO, ROTT, HANS (2004), 'Revision by comparison', *Artificial Intelligence*, 157: 5–47.
- GÄRDENFORS, PETER, MAKINSON, DAVID (1988), 'Revisions of knowledge systems using epistemic entrenchment', in Vardi, M. (ed.), *Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann, Los Altos, CA, 1988, pp. 83–95.
- GROVE, ADAM (1988), 'Two modellings for theory change', *Journal of Philosophical Logic*, 17: 157–170.
- HANSSON, SVEN O. (1999), *A Textbook of Belief Dynamics. Theory Change and Database Updating*, Kluwer, Dordrecht.
- LEVI, ISAAC (2004), *Mild Contraction: Evaluating Loss of Information due to Loss of Belief*, Oxford University Press, Oxford.
- LEWIS, DAVID (1973), *Counterfactuals*, Blackwell, Oxford.
- MEYER, THOMAS, GHOSE, ADITYA, CHOPRA, SAMIR (2001), 'Syntactic representations of semantic merging operations', in *Proceedings of the IJCAI-2001 Workshop on Inconsistency in Data and Knowledge*, Seattle, USA, August 2001, pp. 36–42.
- NAYAK, ABHAYA C. (1994), 'Iterated belief change based on epistemic entrenchment', *Erkenntnis*, 41: 353–390.
- NAYAK, ABHAYA C., PAGNUCCO, MAURICE, PEPPAS, PAVLOS (2003), 'Dynamic belief revision operators', *Artificial Intelligence*, 146: 193–228.
- NAYAK, ABHAYA, GOEBEL, RANDY, ORGUN, MEHMET (2007), 'Iterated belief contraction from first principles', *International Joint Conference on Artificial Intelligence (IJCAI'07)*, pp. 2568–2573.
- NEBEL, BERNHARD (1992), 'Syntax-based approaches to belief revision', in: Gärdenfors, Peter (ed.), *Belief Revision*, Cambridge University Press, Cambridge, pp. 52–88.
- PAGNUCCO, MAURICE, ROTT, HANS (1999), 'Severe withdrawal—and recovery', *Journal of Philosophical Logic*, 28: 501–547. (Full corrected reprint in the JPL issue of February 2000.)
- PAPINI, ODILE (2001), 'Iterated revision operations stemming from the history of an agent's observations', in Williams, Rott (eds.), pp. 279–301.

- PEIRCE, CHARLES S. (1903), 'The nature of meaning', Harvard Lecture delivered on 7 May 1903, published in *The Essential Peirce*, vol. 2 (1803–1913), ed. by the Peirce Edition Project, Indiana University Press, Bloomington, 1998, pp. 208–225.
- RESCHER, NICHOLAS (1964), *Hypothetical Reasoning*, North-Holland, Amsterdam.
- ROTT, HANS (1991a), 'Two methods of constructing contractions and revisions of knowledge systems', *Journal of Philosophical Logic*, 20: 149–173.
- ROTT, HANS (1991b), 'A non-monotonic conditional logic for belief revision I', in Fuhrmann, A., Morreau, M., *The Logic of Theory Change*, LNCS vol. 465, Springer, Berlin, pp. 135–181.
- ROTT, HANS (1992), 'Modellings for belief change: Prioritization and entrenchment', *Theoria*, 58: 21–57.
- ROTT, HANS (2000), '“Just because”: Taking belief bases seriously', in Buss, Samuel R., Hájek, Petr, Pudlák, Pavel (eds.), *Logic Colloquium '98—Proceedings of the 1998 ASL European Summer Meeting*, Lecture Notes in Logic, vol. 13, Urbana, Ill. Association for Symbolic Logic, pp. 387–408.
- ROTT, HANS (2001), *Change, Choice and Inference*, Oxford University Press, Oxford.
- ROTT, HANS (2003), 'Coherence and conservatism in the dynamics of belief. Part II: Iterated belief change without dispositional coherence', *Journal of Logic and Computation*, 13: 111–145.
- ROTT, HANS (2004), 'Stability, strength and sensitivity: converting belief into knowledge', in Brendel, Elke, Jäger, Christoph (eds.), *Contextualisms in Epistemology*, special issue of *Erkenntnis* 61: 469–493.
- ROTT, HANS (2006), 'Revision by comparison as a unifying framework: Severe withdrawal, irrevocable revision and irrefutable revision', *Theoretical Computer Science*, 355: 228–242.
- ROTT, HANS (2007), 'Bounded revision: Two-dimensional belief change between conservatism and moderation', in Rønnow-Rasmussen, Toni et al. (eds.), *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*, internet publication, <http://www.fil.lu.se/hommageawlodek/site/abstra.htm>.
- SEGERBERG, KRISTER (1998), 'Irrevocable belief revision in dynamic doxastic logic', *Notre Dame Journal of Formal Logic*, 39: 287–306.
- SPOHN, WOLFGANG (1988), 'Ordinal conditional functions', in Harper, W.L., Skyrms, B. (eds.), *Causation in Decision, Belief Change, and Statistics, vol. II*, Reidel, Dordrecht, pp. 105–134.
- STALNAKER, ROBERT (1996), 'Knowledge, belief, and counterfactual reasoning in games', *Economics and Philosophy*, 12: 133–163.
- WILLIAMS, MARY-ANNE (1994), 'On the logic of theory base change', in MacNish, C., Pearce, D., Pereira, L.M. (eds.), *Logics in Artificial Intelligence*, LNCS, vol. 838, Springer, Berlin, pp. 86–105.
- WILLIAMS, MARY-ANNE (1995), 'Iterated theory base change: A computational model', in *IJCAI'95—Proceedings of the 14th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo, pp. 1541–1550.
- WILLIAMS, MARY-ANNE, ROTT, HANS (eds.), (2001), *Frontiers in Belief Revision*, Kluwer, Dordrecht.

Appendix: Sphere pictures




Figure 1: Conservative expansion




Figure 5: Very radical expansion




Figure 2: Plain expansion




Figure 6: Radical revision (= 'irrevocable revision')




Figure 3: Moderate expansion




Figure 7: Very radical revision




Figure 4: Radical expansion




Figure 8: Conservative revision (= 'natural revision')




Figure 9: Moderate revision (= 'lexicographic revision')




Figure 13: Plain severe revision




Figure 10: Restrained revision




Figure 14: Moderate severe revision




Figure 11: Severe withdrawal (= 'mild contraction')




Figure 15: Conservative contraction (\approx AGM contraction)




Figure 12: Severe revision




Figure 16: Moderate contraction




Figure 17: Refining (= 'Reverse lexicographic belief change')




Figure 21: Dual to severe withdrawal




Figure 18: Revision by comparison (= 'raising')




Figure 22: Gentle raising




Figure 19: Irrefutable revision (with fixed ϵ)




Figure 23: Gentle lowering




Figure 20: Lowering




Figure 24: Raising by strict comparison




Figure 25: Lowering by strict comparison




Figure 27: Bounded revision (non-strict version)




Figure 26: Bounded revision (strict version)

HANS ROTT
 Institut für Philosophie
 Universität Regensburg
 93040 Regensburg, Germany
 Hans.Rott@psk.uni-regensburg.de