TWO-DIMENSIONAL BELIEF CHANGE
An Advertisement

Hans Rott

Department of Philosophy, University of Regensburg
hans.rott@psk.uni-r.de

Abstract

In this paper I compare two different the models of two-dimensional belief change, namely 'revision by comparison' (Fermé and Rott, *Artificial Intelligence* 157, 2004) and 'bounded revision' (Rott, in *Hommage à Wlodek*, Uppsala 2007). These revision operations are two-dimensional in the sense that they take as arguments pairs consisting of an input sentence and a reference sentence. Two-dimensional revision operations add a lot to the expressive power of traditional qualitative approaches to belief revision and refrain from assuming numbers as measures of degrees of belief.

1. Introduction

Representations of belief states in terms of probability functions or ranking functions are very rich and powerful. However, it is often hard to come by meaningful numbers. Qualitative belief change in the style of Alchourrón, Gärdenfors and Makinson (1985, henceforth 'AGM') and its extensions to iterated belief change in the 1990s, on the other hand, are simple and do not need numbers, but are a lot more restricted in their expressive power. Two-dimensional belief revision attempts to strike a good balance between the advantages of AGM-style qualitative and quantitative approaches.

1.1. Revision by comparison

Fermé and Rott (2004) suggested a basically qualitative approach that is more flexible than AGM style models in that it allows a new piece of information to be accepted in various degrees or strengths. The key idea is that an input sentence

Dagstuhl Seminar Proceedings 07351
Formal Models of Belief Change in Rational Agents
http://drops.dagstuhl.de/opus/volltexte/2007/1240

α does not come with a number, but does not come 'naked' either. It rather comes with a reference sentence δ that typically expresses an antecedently held belief. The agent is then supposed to follow an instruction of the form

'Accept α with a strength that at least equals that of δ.'

Revision by comparison is a model that lies between the traditional qualitative and quantitative approaches.[1] In the context of revision by comparison, 'revising by α' is understood to mean revising with reference to some existing belief expressed by the reference sentence δ.

A drawback of the revision-by-comparison approach of Fermé and Rott is that it does not satisfy the Darwiche-Pearl postulates for iterated belief change (Darwiche and Pearl 1997). These postulates have a very appealing possible worlds semantics that strongly suggests that they should be satisfied by iterated revision functions (compare Section 2.2 below).

1.2. Bounded revision

Bounded revision was motivated by the same concerns as revision by comparison, combined with the desire to satisfy the Darwiche-Pearl postulates (Rott 2007). Intuitively, δ serves as a bound for the acceptance of α. The reference sentence δ functions here as a measure of how firmly entrenched α should be in the agent's belief state after the change.

The idea of bounded revision can be expressed more precisely by the following recipe:

'Accept α as long as δ holds along with α, and just a little more.'

In a way, the acceptance of α is *bounded* by δ. We can think of the reference sentence δ in two ways. First, we may suppose that it is a marker delineating the shape of a sphere in a Grovean system of spheres that characterizes the reasoner's initial belief state. Second, we can use δ as a sentence that is supposed to hold in a range of relatively plausible situations in which α holds. Rott (2007) argues that the latter option is preferable for bounded revision. One advantage is that the second option is more general because it includes the first option as a special case.

[1] An approach similar to revision by comparison was introduced earlier in Cantwell's (1997) 'raising' operation. Cantwell has also presented an interesting dual operation that he calls 'lowering'.

Any arbitrary sentence δ may sensibly serve as the parameter sentence for a bounded revision. The agent need not actually believe that δ is true. However, the paradigm cases are those in which δ is cotenable with α to some extent, in the sense that a stretch of the comparatively plausible ways of making α true are all ways that make δ true as well.[2] The greater the stretch where δ holds along with α, the firmer α gets accepted by a revision that is bounded by δ.

Usually the intended cases of belief revision are those in which the input sentence α is not believed prior to the revision. However, two-dimensional revision may well be used to increase the strength or entrenchment of a sentence that the agent has already believed to be true prior to the revision.[3]

2. Generalizing AGM to two-dimensional revisions of belief states

What is a belief state? For the purposes of this paper, belief states may be entities of any type whatsoever, neural states, holistic mental states, abstract machine states etc. We assume that the set of beliefs of a reasoner, but not necessarily his or her whole belief state is epistemically accessible. The beliefs are, so to speak, the visible tip of the iceberg that itself remains concealed from our eyes and, perhaps, from the reasoner's own eyes as well. Our hypothesis is that belief states have a rich structure that determines at least the development of the agent's belief set in response to any sequence of inputs. We do not exclude that it determines more. In Section 2.3, we shall specify concrete formal structures as representations of belief states that contain a lot more structure than a plain belief set, but are still abstractions from 'real' belief states. We start this section, however, by addressing the problem of belief *change* in abstraction from any particular conceptualization of belief states.

2.1. One-dimensional and two-dimensional belief revision operators

A *one-dimensional* belief revision operation is a function $*$ that takes a belief state $\mathcal{B}$ and an input sentence α and returns the new belief state. $*(\mathcal{B},\alpha)$ denotes the state $\mathcal{B}$ revised by α. A *two-dimensional* revision function is similar, except that the input is a pair of sentences $\langle\alpha,\delta\rangle$. The first sentence is the *input*

[2] Intuitively, not only the *most* plausible α-worlds, but also all those that are *sufficiently* plausible are moved center stage in a bounded revision by α, and it is precisely the task of the reference sentence δ to characterize what is meant 'sufficient.' Compare Fig. 1 below.

[3] This is the main idea underlying Cantwell's (1997) approach.

sentence, the second sentence the *reference sentence*.[4] The sentences can vary fully independently from each other, thus the name 'two-dimensional.'[5] Usually, I will use the variables α, β and γ etc. for input sentences, and the variables δ, ε, ζ etc. for reference sentences.

We work with a finite propositional language $\mathcal{L}$. The set of possible worlds (interpretations, models) and the set of sets of logically equivalent sentences are then finite, too.[6] We use Cn to indicate a consequence operation governing $\mathcal{L}$. We suppose throughout this paper that the logic is Tarskian, that it includes classical propositional logic, and that it satisfies the deduction theorem.[7] The only inconsistent and logically closed set in the language is the set of all sentences which we also denote by $\mathcal{L}$.

Notation: If Γ is a set of sentences and α and β are sentences, I shall write $\Gamma + \alpha$ for $\Gamma \cup \{\alpha\}$. For any belief state $\mathcal{B}$, $\ulcorner\mathcal{B}\urcorner$ is the set of beliefs held by a person in belief state $\mathcal{B}$ (more exactly: the beliefs that can be ascribed to the person, or the beliefs that the person is committed to). We assume that $\ulcorner\mathcal{B}\urcorner$ is logically closed. If $\mathcal{B}$ and $\mathcal{B}'$ are two belief states, then $\mathcal{B} \simeq \mathcal{B}'$ is short for $\ulcorner\mathcal{B}\urcorner = \ulcorner\mathcal{B}'\urcorner$. As is common in theories of one-dimensional revision functions $*$, we write $\mathcal{B} * \alpha$ for $*(\mathcal{B}, \alpha)$. For a two-dimensional revision function $*$, we write $\mathcal{B} *_\delta \alpha$ for $*(\mathcal{B}, \langle\alpha, \delta\rangle)$.

The main benchmark in theory change are the famous AGM postulates for one-step belief revision (AGM 1985). We rewrite them in a new notation that makes explicit that belief revision, (i), is really about the revision of belief states rather than belief sets and, (ii), is more generally conceived as two-dimensional rather than one-dimensional.

---

[4]The terminology of input and reference sentences is taken over from Fermé and Rott (2004).

[5]The overused epithet 'two-dimensional' is certainly not ideal. Unfortunately, all the friendly suggestions of colleagues (thank you!) seem to have other drawbacks, so I stick to my name as long as no better one springs to mind.

[6]We presuppose finiteness mainly as a matter of convenience, in order not to burden this paper with technical details distracting us from the main issues. An infinite language would not complicate things as long as we work with entrenchment relations, but when working with systems of spheres, infinity complicates the matter enormously. See, e.g., Pagnucco and Rott (1999, Section 8).

[7]By saying that the logic Cn is Tarskian, we mean that it is reflexive ($\Gamma \subseteq Cn(\Gamma)$), monotonic (if $\Gamma \subseteq \Gamma'$, then $Cn(\Gamma) \subseteq Cn(\Gamma')$), idempotent ($Cn(Cn(\Gamma)) \subseteq Cn(\Gamma)$) and compact (if $\alpha \in Cn(\Gamma)$, then $\alpha \in Cn(\Gamma')$ for some finite $\Gamma' \subseteq \Gamma$). The deduction theorem says that $\alpha \rightarrow \beta \in Cn(\Gamma)$ if and only if $\beta \in Cn(\Gamma \cup \{\alpha\})$. We may write $\Gamma \vdash \alpha$ for $\alpha \in Cn(\Gamma)$.

(AGM1) $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is logically closed.

(AGM2) $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ contains α

(AGM3) $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is a subset of $Cn(\ulcorner \mathcal{B} \urcorner + \alpha)$

(AGM4) If α is consistent with $\ulcorner \mathcal{B} \urcorner$, then $\ulcorner \mathcal{B} \urcorner$ is a subset of $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$

(AGM5) If α is consistent, then $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is consistent

(AGM6) If α is logically equivalent with β, then $\ulcorner \mathcal{B} *_\delta \alpha \urcorner = \ulcorner \mathcal{B} *_\delta \beta \urcorner$

(AGM7) $\ulcorner \mathcal{B} *_\delta (\alpha \wedge \beta) \urcorner$ is a subset of $Cn(\ulcorner \mathcal{B} *_\delta \alpha \urcorner + \beta)$

(AGM8) If β is consistent with $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$, then $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is a subset of $\ulcorner \mathcal{B} *_\delta (\alpha \wedge \beta) \urcorner$

To get the AGM postulates for one-dimensional revision operations, just drop the subscript 'δ' from each occurrence of '$*_\delta$'.

It follows from (AGM3) and (AGM4) that if $\ulcorner \mathcal{B} \urcorner$ is consistent, then $\ulcorner \mathcal{B} \urcorner = \ulcorner \mathcal{B} *_\delta \top \urcorner$ for any δ. The beliefs in a consistent belief state $\mathcal{B}$ are exactly the same beliefs as after the revision of $\mathcal{B}$ by the tautology $\top$, regardless of which reference sentence δ is given.

From the point of view of the present paper, (AGM5) introduces an unnecessary loss of generality.[8] An agent may consider more sentences than just logical falsehoods as 'absolutely impossible'. And a revision by an absolutely impossible sentence may lead him or her either into an inconsistent belief set (the AGM idea) or into a refusal to change anything (an alternative idea which makes equally good sense).

It is natural to assume that two-dimensional revision functions satisfy the AGM postulates, except for (AGM5). But as we shall soon see, revision by comparison is an operation that does not satisfy them. The reason is that it has features of belief contraction (belief withdrawal, removal, subtraction, etc.) as well as belief revision. This does not hold for bounded revision which is thus more well-behaved or simpler in this respect.

(AGM6)–(AGM8) compare revisions by two different input sentences. It seems absolutely unproblematic to validly strengthen each of these postulates by replacing the last occurrence of 'δ' by a *logically equivalent* reference sentence 'ε'. A more interesting and more important question relevant for two-dimensional belief change is whether it is legitimate to vary the reference sentence attached to '$*$' within a single postulate, and use an *arbitrary*, not necessarily equivalent

[8] Recommendations how to weaken (AGM5) in a one-dimensional context are given in Rott (2001, pp. 149–153, 206 and 118).

sentence 'ε' in the place of the second occurrence of 'δ'. This seems plausible insofar as we can assume that the reference sentence only specifies the strength with which the input sentence is accepted, but not the content of the new belief set. In other words, insofar as the beliefs obtained after a revision by an input sentence α do not depend on the reference sentence δ at all. A central condition to be discussed is this:

(SBC) $\ulcorner \mathcal{B} *_\delta \alpha \urcorner = \ulcorner \mathcal{B} *_\varepsilon \alpha \urcorner$ for all δ and ε

Let us call this condition the *Same Beliefs Condition*. It says that the posterior *belief set* does not depend on the reference sentence. If (SBC) is satisfied, only the structure of the *belief state* obtained is sensitive to variations of the reference sentence. We shall see that bounded revision satisfies (SBC), but revision by comparison in general does not. In the latter model, the relative strengths of input and reference sentence do matter, which is again due to the fact that revision by comparison is not a pure operation of revision but has features of contraction, too. Restricted versions of (SBC), however, will turn out to hold for revision by comparison as well.

2.2. Iterations

Now we turn to the *Darwiche-Pearl postulates* for iterated belief change that were already mentioned in the introduction. We adapt them here so as to apply to the general case of two-dimensional revision functions.

(DP1) If β implies α, then $(\mathcal{B} *_\delta \alpha) *_\varepsilon \beta \simeq \mathcal{B} *_\zeta \beta$

(DP2) If β is inconsistent with α, then $(\mathcal{B} *_\delta \alpha) *_\varepsilon \beta \simeq \mathcal{B} *_\zeta \beta$

(DP3) If α is in $\ulcorner \mathcal{B} *_\zeta \beta \urcorner$, then α is in $\ulcorner (\mathcal{B} *_\delta \alpha) *_\varepsilon \beta \urcorner$

(DP4) If $\neg\alpha$ is not in $\ulcorner \mathcal{B} *_\zeta \beta \urcorner$, then $\neg\alpha$ is not in $\ulcorner (\mathcal{B} *_\delta \alpha) *_\varepsilon \beta \urcorner$

A much more cautious formulation would take the same reference sentence 'δ' everywhere rather than varying the reference sentences within each condition. It remains to be seen how various two-dimensional revision operations fare with respect to the more cautious and the bolder formulations of (DP1)–(DP4). The two approaches considered in this paper are clear cases anyhow: Bounded revision satisfies, but revision by comparison does not satisfy the Darwiche-Pearl postulates.

Since these postulates concern iterations, they make statements not only about *belief sets*, but implicitly also about one-step changes of *belief states* as well.

As Darwiche and Pearl (1997) have shown, (the one-dimensional counterparts of) these postulates correspond one by one to very appealing semantic conditions in terms of a total preordering of possible worlds (graphically representable as a Grovean system of spheres.[9] Assume that a belief state is represented by such a preordering and that the new piece of information is α. Then (DP1) and (DP2) essentially say that a revision by α should preserve the preordering *restricted to the α-worlds* and the preordering *restricted to the $\neg\alpha$-worlds*, respectively. (DP3) and (DP4) taken together essentially say that the relative position of an α-world with respect to a $\neg\alpha$-world must not be worse after a revision of the belief state by α. This eminently plausible semantics recommends that the Darwiche-Pearl postulates be obeyed by reasonable iterated belief revision operators, and in fact it can be shown that an important class of iterated revision functions obey them, amongst them bounded revision (Rott 2007).

2.3. Representing belief states as order relations

In what sense do equations like the above characterize an iterated revision function for belief states? How can we get from $\mathcal{B}$ and an input of the form α or $\langle\alpha, \beta\rangle$ to the revised belief state? A key to understanding much of traditional research in belief revision is that the *belief sets* obtained after potential second revision steps give sufficient evidence about the *belief state* the agent is in after the first revision step.

We will work with two different forms of representation of belief states that are sufficient to determine the set of beliefs held after any sequence of inputs.[10] The first is a total pre-ordering of possible worlds. Such preorderings can equivalently be presented graphically in the form of Grovean systems of spheres (s.o.s.) of possible worlds (Grove 1988). This is by far the most easily comprehensible representation. For this paper, we assume that an s.o.s. $\$$ is a non-empty, finite set of sets of possible worlds such that for any two sets S and S' in $\$$, either $S \subseteq S'$ or $S' \subseteq S$ (that is, the elements of $\$$ are 'nested', or form a chain with respect to $\subseteq$). Intuitively, the most plausible worlds are contained in the smallest (graphically, 'innermost') sphere of $\$$, the second most plausible worlds are contained in the second smallest sphere, and so on. Worlds not contained in any sphere are called inaccessible in $\$$. The set of sentences true at all the worlds contained in the innermost sphere express the beliefs held true by an

[9] Grove (1988). More on this in Section 2.3.

[10] There is a third representation in terms of prioritized belief bases which is particularly attractive for the operation of revision by comparison; cf. Rott (forthcoming).

agent in belief state $\$$; this is the agent's belief set and denoted by $\ulcorner \$ \urcorner$.

Our second way of representing a belief state is by a total ordering $\leq$ of sentences, usually called 'entrenchment relation' (Gärdenfors and Makinson (1988), Rott (2001, 2003a)). Such an ordering can roughly be thought of as reflecting the degree of belief or the comparative retractability of the respective sentences. These degrees are required to respect logical structure in two ways. First, if α implies β, then the entrenchment of α cannot be higher as that of β ('dominance'). Second, the conjunction $\alpha \wedge \beta$ is not less entrenched than the weaker of α and β ('conjunctiveness'). In the first respect entrenchments behave like probabilities, in the second, they are quite different. The set of sentences that are more than minimally entrenched are the beliefs held true by an agent in belief state $\leq$; this is the agent's belief set and denoted by $\ulcorner \leq \urcorner$.[11]

What does it mean to say that a system of spheres or an entrenchment ordering represents a belief state? This is not a trivial question. As we said before, a formal structure like an s.o.s. $\$$ or an entrenchment relation $\leq$ is still an abstraction from a 'real' belief state. And what the real belief state is may be inscrutable to us. But we can say that $\$$ or $\leq$ *represents a belief state* if it reproduces just that aspect of belief states we may hope to have access to, i.e., the development of the agent's beliefs. The structures $\$$ and $\leq$ should encode all the information necessary to derive the resulting belief sets for all iterated belief changes, provided that a specific method of using $\$$ and $\leq$ to construct a single revision step is given. For instance, in one-dimensional belief revision this means that for systems of spheres

$$\ulcorner (((\mathcal{B} * \alpha) * \beta) * \gamma) * \ldots \urcorner = \ulcorner (((\$^*_\alpha)^*_\beta)^*_\gamma)^* \ldots \urcorner$$

and in two-dimensional belief revision

$$\ulcorner (((\mathcal{B} *_\delta \alpha) *_\varepsilon \beta) *_\zeta \gamma) * \ldots \urcorner = \ulcorner (((\$^*_{\alpha,\delta})^*_{\beta,\varepsilon})^*_{\gamma,\zeta})^* \ldots \urcorner$$

for all finite sequences of inputs $\langle \alpha, \beta, \gamma, \ldots \rangle$ or $\langle \langle \alpha, \delta \rangle, \langle \beta, \varepsilon \rangle, \langle \gamma, \zeta \rangle, \ldots \rangle$, respectively. The definition for entrenchment relations is similar.

It is well-known from the belief revision literature beginning with AGM that an ordering representation of a belief state $\mathcal{B}$ determines a one-dimensional revision function specifying, for each potential input sentence α, the belief set that results from revising $\mathcal{B}$ by α. Conversely, given such a one-dimensional revision

[11] We assume in this paper that the entrenchment relation is non-trivial in the sense that $\top$ is more entrenched than $\bot$.

function satisfying certain 'rationality postulates', one can (re-)construct an ordering that can be taken to underlie that revision function.[12]

For the connection between systems of spheres and revised belief sets[13], we can make use of the following transitions (see Grove 1988):

(From \$ to $\ulcorner * \urcorner$) β is in $\ulcorner \mathcal{B} * \alpha \urcorner$ if and only if there is a sphere in \$ containing some α-worlds which are all β-worlds, or there is no sphere containing any α-worlds.

(From $\ulcorner * \urcorner$ to \$) A set S of possible worlds is in \$ if and only if there is a sentence α such that $S = \{w \in W :$ for some β, w satisfies all sentences in $\ulcorner \mathcal{B} * (\alpha \vee \beta) \urcorner\}$.

For the connection between entrenchments and revised belief sets, we can use the following transitions (see Gärdenfors and Makinson 1988, Rott 2001, Ch. 8):

(From $\leq$ to $\ulcorner * \urcorner$) β is in $\ulcorner \mathcal{B} * \alpha \urcorner$ if and only if $\neg\alpha < \alpha \rightarrow \beta$ or $\top \leq \neg\alpha$.

(From $\ulcorner * \urcorner$ to $\leq$) $\alpha \leq \beta$ if and only if α is not in $\ulcorner \mathcal{B} * \neg(\alpha \wedge \beta) \urcorner$ or $\ulcorner \mathcal{B} * \neg(\alpha \wedge \beta) \urcorner$ is inconsistent.

There is a certain asymmetry in the intuitive plausibilities of these recipes. For system of spheres, the construction of $*$ in terms of \$ is more transparent then the (re-)construction of \$ from $*$. But for entrenchments, the (re-)construction of $\leq$ from $*$ is more convincing than the construction of $*$ in terms of $\leq$. Because the two directions fit together perfectly, for systems of spheres as well as for entrenchment relations, both modellings are almost universally accepted in the one-dimensional setting.

Additional support comes from the result that the s.o.s. modelling and the entrenchment modelling are equivalent in quite a strong sense. One can easily complete the circle of equivalences, by defining an entrenchment relation from an s.o.s. and an s.o.s. from an entrenchment relation, that are equivalent in the sense that they generate exactly the same revision function. The relevant transitions are (see, for instance, Pagnucco and Rott 1999):

(From \$ to $\leq$) $\alpha \leq \beta$ if and only if for all S in \$, if α is true throughout S, then β is true throughout S as well.

[12]The additional information encoded in two-dimensional belief change operations is not needed for the (re-)construction of the belief state.

[13]Note that the following connections, as well as the corresponding ones connecting revised belief sets and entrenchments, do not appeal to the belief *states*.

(From $\leq$ to \$) A non-empty set S of possible worlds is in \$ if and only if there is a sentence α such that $S = \{w \in W : w$ satisfies all sentences β with $\alpha \leq \beta\}$.[14]

As long as there is no danger of confusion, we will allow ourselves to say that an s.o.s. or an entrenchment relation *is* a belief state rather than saying that it is an abstraction from, or a representation of, a belief state.

3. Bounded revision

Bounded revision is a change operation the non-iterated part of which is fully taken care of by the AGM postulates (with the above-mentioned exception of AGM5). So we can immediately address the problem of iterations.

3.1. Bounded revision as an operation for iterated belief change

The general iteration condition for the two-dimensional operation of *bounded revision* is this:

(IT^b)

$$\mathcal{B} *_\delta \alpha *_\varepsilon \beta \;\simeq\; \begin{cases} \mathcal{B} *_\zeta (\alpha \wedge \beta) & \text{if } \ulcorner \mathcal{B} *_\delta (\alpha \wedge (\delta \rightarrow \beta)) \urcorner \text{ is consistent with } \beta \\ \mathcal{B} *_\zeta \beta & \text{otherwise} \end{cases}$$

Unfortunately, this condition is not very transparent. The rationale for it will become clear when we turn to the modellings in terms of systems of spheres and entrenchments.

Look at what the condition of bounded revision gives for the important special case $\beta = \top$. We obtain, after a little simplification using (AGM6),

$$\mathcal{B} *_\delta \alpha *_\varepsilon \top \;\simeq\; \begin{cases} \mathcal{B} *_\zeta \alpha & \text{if } \ulcorner \mathcal{B} *_\delta \alpha \urcorner \text{ is consistent} \\ \mathcal{B} *_\zeta \top & \text{otherwise} \end{cases}$$

If $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is consistent, the left hand side equals $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$, by (AGM3) and (AGM4). Therefore, (IT^b) implies that if $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is consistent, it is identical with $\ulcorner \mathcal{B} *_\zeta \alpha \urcorner$ for any reference sentence ζ. By a symmetrical argument applied $\ulcorner \mathcal{B} *_\zeta \alpha \urcorner$, it also follows that if $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is inconsistent so is $\ulcorner \mathcal{B} *_\zeta \alpha \urcorner$. In sum, then, bounded revision satisfies the Same Beliefs Condition (SBC). Knowing this, we

[14] I neglect the problem of adding the empty set to systems of spheres. Cf. footnote **??**.

will drop the subscripts to '$*$' in this section when we are only interested in the belief state resulting from bounded revision.

We consider two limiting cases of reference sentences δ that are (i) never or (ii) always cotenable with α. The operations of conservative revision and moderate revision[15] can be obtained from bounded revision by fixing the parameter sentence as $\bot$ (FALSITY) and $\top$ (TRUTH), respectively. Both conservative and moderate revision, however, seem to be defective. Conservative revision accepts new evidence, but always accords the lowest possible entrenchment to it, so that it gets immediately lost if a contradiction with another piece of evidence arises (see Rott 2003b). Moderate revision in a way suffers from the opposite defect by accepting the new information very firmly. While conservative revision is too conservative, moderate revision seems too radical. Bounded revision covers the range between these two extremes.

For the first limiting case, let δ be $\bot$ (or equivalently, let δ be $\neg\alpha$). Then (IT^b) reduces to

$$\mathcal{B} *_\bot \alpha * \beta \;\simeq\; \begin{cases} \mathcal{B} * (\alpha \wedge \beta) & \text{if } \ulcorner \mathcal{B} * \alpha \urcorner \text{ is consistent with } \beta \\ \mathcal{B} * \beta & \text{otherwise} \end{cases}$$

which characterizes *conservative revision*. The upper line follows already from the AGM postulates (AGM7) and (AGM8) for one-step revisions.

For the second limiting case, let δ be $\top$ (or equivalently, let δ be α). Then, given the AGM postulates minus (AGM5), (IT^b) reduces to

$$\mathcal{B} *_\top \alpha * \beta \;\simeq\; \begin{cases} \mathcal{B} * (\alpha \wedge \beta) & \text{if } \ulcorner \mathcal{B} * (\alpha \wedge \beta) \urcorner \text{ is consistent} \\ \mathcal{B} * \beta & \text{otherwise} \end{cases}$$

which characterizes *moderate revision*. Had we also accepted the consistency postulate (AGM5), the upper line could be made conditional on the simple requirement that $\alpha \wedge \beta$ be consistent.

The revision operation specified by the AGM axioms (except (AGM5)) and the iteration axiom IT^b) satisfies the Darwiche-Pearl postulates. This can either be proved directly on the level of the postulates, or by considering the s.o.s.

[15]These are my names (Rott 2003b). The operations I denote by these names are more well-known as *natural revision* as introduced by Boutilier (1993) and *lexicographic revision* as studied by Nayak (1994) and others, respectively.

semantics of bounded revision presented in the next section.[16]

3.2. Bounded revision as an operation on systems of spheres

In Section 2.3, we have seen that given the revised belief sets $\ulcorner \mathcal{B} * \alpha \urcorner$ for all potential inputs α, one can reconstruct an ordering representation of the belief state $\mathcal{B}$. This throws new light on the equation (IT^b) for iterated belief change from Section 3.1. Given the revised belief sets $\ulcorner (\mathcal{B} *_\delta \alpha) * \beta \urcorner$ for all potential inputs β, one can similarly reconstruct an ordering representation of the belief state $\mathcal{B} *_\delta \alpha$. That is to say that (IT^b) in effect specifies a transition from a representation of the belief state $\mathcal{B}$ to a representation of the belief state $\mathcal{B} *_\delta \alpha$.

Bounded revision functions should thus be viewed as functions applying to formal representations of belief states (not only on belief sets which contain too little information, and not on belief states which may ultimately be inscrutable). In this and the next subsection, we give a direct account of the relevant transitions, as applying on systems of spheres and entrenchment relations respectively. We begin with the representation of belief states in terms of systems of spheres.

Let $[\alpha]$ denote the sets of possible worlds in which α is true. Let α intersect \$, i.e., let there be at least one sphere in \$ that has a non-empty intersection with $[\alpha]$. Let $S_{\alpha,\delta}$ be the smallest sphere S in \$ such that $S \cap [\alpha] \not\subseteq [\delta]$; if there is no such sphere, take $S_{\alpha,\delta}$ to be the largest sphere in \$.

Let $\$^*_{\alpha,\delta}$ denote the system of spheres that results from revising the prior s.o.s. \$ by an input sentence α, bounded by reference sentence δ. The following definition of bounded revision as an operation on systems of spheres applies to the case in which α intersects \$.

(BoundRevSS)

$$\$^*_{\alpha,\delta} = \{S \cap [\alpha] : S \in \$, S \cap [\alpha] \neq \emptyset \text{ and } S \subseteq S_{\alpha,\delta}\} \cup \{S \cup ([\alpha] \cap S_{\alpha,\delta}) : S \in \$\}$$

If α does not intersect \$, then we simply put $\$^*_{\alpha,\delta} = \$ \cup \{\emptyset\}$. We assume that once a world is inaccessible, it cannot be made accessible by a revision operation.[17]

[16] A formulation of the systems of spheres semantics for the Darwiche-Pearl postulates is given in Rott (2007).

[17] This question is particularly relevant for 'moderate' belief change.

Figure 1 shows what happens to an s.o.s. when it gets revised by bounded revision. The numbers used in this figure are, of course, just there to indicate the relative plausibilities of (regions of) possible worlds. '1' designates the most plausible worlds, '2' the second most plausible worlds, and so on. '∞' designates the doxastically impossible or inaccessible worlds.

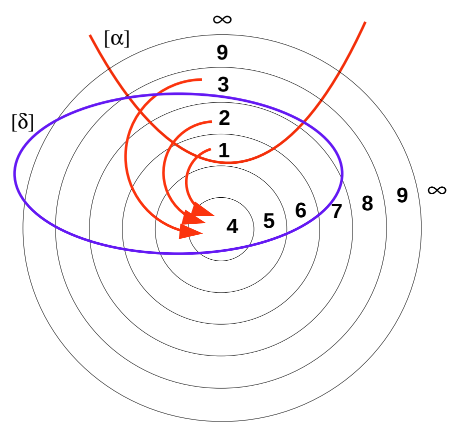

Fig. 1: Bounded revision: Moving the best $\alpha \wedge \delta$-worlds (and a few more α-worlds) to the center, preserving all distinctions

The picture brings out the fact that bounded revision tends to increase the number of spheres in an s.o.s. (in the example from 6 to 9), thus making plausibility distinctions finer.

Now let us have a look at the limiting cases for (BoundRevSS).[18] If δ is $\bot$ (or $\neg\alpha$), then $S_{\alpha,\delta}$ is the smallest sphere S in \$ such that $S \cap [\alpha] \neq \emptyset$; let us denote this sphere by S_α. Then we get conservative revision:

$$\$^*_{\alpha,\bot} \;=\; \{S_\alpha \cap [\alpha]\} \;\cup\; \{S \cup (S_\alpha \cap [\alpha]) : S \in \$\}$$

If δ is $\top$ (or α), then $S_{\alpha,\delta}$ is the largest sphere S in \$; let us denote this sphere by S_{max}. Then we get moderate revision:

$$\$^*_{\alpha,\top} \;=\; \{S \cap [\alpha] : S \in \$ \text{ and } S \cap [\alpha] \neq \emptyset\} \;\cup\; \{S \cup (S_{max} \cap [\alpha]) : S \in \$\}$$

Assuming that belief states can be represented by systems of spheres, we get the following characterization theorem:

[18] I neglect the case of an impossible α now.

Observation 1. (i) The two-dimensional revision function $*$ determined by (BoundRevSS) satisfies (AGM1)–(AGM8) and (IT^b).

(ii) If the two-dimensional revision function $*$ satisfies (AGM1)–(AGM8) and (IT^b), then there is, for each iterated revision process determined by $*$, a system of spheres \$ such that at each state in this process, the set of beliefs accepted is identical with the set of beliefs determined by the corresponding system of spheres as transformed according to (BoundRevSS):

$$\ulcorner \mathcal{B} *_\delta \alpha \urcorner = \ulcorner \$^*_{\alpha,\delta} \urcorner$$
$$\ulcorner (\mathcal{B} *_\delta \alpha) *_\varepsilon \beta \urcorner = \ulcorner (\$^*_{\alpha,\delta})^*_{\beta,\varepsilon} \urcorner$$
$$\ulcorner ((\mathcal{B} *_\delta \alpha) *_\varepsilon \beta) *_\zeta \gamma \urcorner = \ulcorner ((\$^*_{\alpha,\delta})^*_{\beta,\varepsilon})^*_{\gamma,\zeta} \urcorner$$

... and so on.

The proof of this observation has to be supplied elsewhere.

3.3. Bounded revision as an operation on entrenchment relations

We now turn to the direct account of transitions from representations of $\mathcal{B}$ to representations of $\mathcal{B} * \alpha$ in terms of entrenchment relations. Bounded revision thus becomes an iterable revision operation operating on entrenchment relations.

Let $\leq^*_{\alpha,\delta}$ denote the entrenchment relation that results from revising the prior entrenchment relation $\leq$ by an input sentence α, bounded by reference sentence δ. The definition of bounded revision as an operation on entrenchment relations applies (at least) for cases in which $\neg\alpha$ is less entrenched than the tautology $\top$.[19] Here is the ordering, defined by comparing any two sentences β and γ:

(BoundRevEnt)

$$\beta \leq^*_{\alpha,\delta} \gamma \text{ iff } \begin{cases} \alpha \rightarrow \beta \leq \alpha \rightarrow \gamma & \text{, if } \alpha \rightarrow (\beta \wedge \gamma) \leq \alpha \rightarrow \delta \\ & \quad \text{and } \alpha \rightarrow (\beta \wedge \gamma) < \top \\ \beta \leq \gamma & \text{, otherwise} \end{cases}$$

It can be proved that (BoundRevEnt) leads from entrenchment relations to

[19] If, on the other hand, $\top \leq \neg\alpha$, then (the lower line of) the following condition rules that the entrenchment relation should not change at all. In order to satisfy the 'success' condition (AGM2) when α is impossible, we would have to stipulate that in this case $\ulcorner \leq \urcorner$ be the inconsistent set $\{\phi : \bot \leq \phi\}$.

entrenchment relations, i.e. transitive relations that satisfy dominance and conjunctiveness.[20]

With (BoundRevEnt), we get the posterior ordering $\delta <^*_{\alpha,\delta} \alpha$, that is $\delta \leq^*_{\alpha,\delta} \alpha$ and not $\alpha \leq^*_{\alpha,\delta} \delta$. But α surpasses δ after the revision only to the slightest possible degree: There is no sentence ϕ for which $\delta <^*_{\alpha,\delta} \phi <^*_{\alpha,\delta} \alpha$. Thus one could say that (BoundRevEnt) defines some kind of 'revision by comparison', in that it implements a reasonable way of minimally accepting the condition $\delta < \alpha$.

As mentioned before, we can allow arbitrary sentences to take the role of δ, but then of course δ cannot be interpreted as specifying a degree of belief relative to the current belief state (characterized by $\leq$). In fact δ need not even be a belief at all in this interpretation. Limiting cases are again obtained when the reference sentence is a logical falsehood or a logical truth. If δ is $\bot$ then (BoundRevEnt) reduces to

$$\beta \leq^*_{\alpha,\bot} \gamma \ \text{ iff } \ \begin{cases} \alpha \rightarrow \beta \leq \alpha \rightarrow \gamma & \text{if } \alpha \rightarrow (\beta \wedge \gamma) \leq \neg\alpha \text{ and } \alpha \rightarrow (\beta \wedge \gamma) < \top \\ \beta \leq \gamma & \text{otherwise} \end{cases}$$

which is the entrenchment recipe for conservative revision.[21] If δ is $\top$ then (BoundRevEnt) reduces to

$$\beta \leq^*_{\alpha,\top} \gamma \ \text{ iff } \ \begin{cases} \alpha \rightarrow \beta \leq \alpha \rightarrow \gamma & \text{if } \alpha \rightarrow (\beta \wedge \gamma) < \top \\ \beta \leq \gamma & \text{otherwise} \end{cases}$$

which is the entrenchment recipe for moderate revision.[22]

Assuming that belief states can be represented by entrenchment relations, we get the following characterization theorem:

Observation 2. (i) The two-dimensional revision function $*$ determined by (BoundRevEnt) satisfies (AGM1)–(AGM8) and (IT^b).

[20] (BoundRevEnt) violates the success condition (AGM2) if $\top \leq \neg\alpha$ and we use the usual definition of $\ulcorner \leq^*_{\alpha,\delta} \urcorner$. The belief set associated with any non-trivial entrenchment relation is consistent. If we insist that a revision by an 'impossible' input results in the inconsistent belief set, we need to redefine $\ulcorner \leq^*_{\alpha,\delta} \urcorner = \{\beta : \bot \leq^*_{\alpha,\delta} \beta\}$ in such a case.

[21] As usual in the AGM paradigm, $\alpha \rightarrow (\beta \wedge \gamma) \leq \neg\alpha$ means $\beta \wedge \gamma \notin K * \alpha$. If this condition is satisfied then $\alpha \rightarrow \beta \leq \alpha \rightarrow \gamma$ reduces to $\alpha \rightarrow \beta \leq \neg\alpha$.

[22] Had we presumed that only logical truths get top entrenchment (an assumption corresponding to (AGM5)), then $\alpha \rightarrow (\beta \wedge \gamma) < \top$ had meant the same as $\beta \wedge \gamma \notin Cn(\alpha)$.

(ii) If the two-dimensional revision function $*$ satisfies (AGM1)–(AGM8) and (IT^b), then there is, for each iterated revision process determined by $*$, an entrenchment relation $\leq$ such that at each state in this process, the set of beliefs accepted is identical with the set of beliefs determined by the corresponding entrenchment relation as transformed according to (BoundRevEnt):

$$\ulcorner \mathcal{B} *_\delta \alpha \urcorner = \ulcorner \leq^*_{\alpha,\delta} \urcorner$$
$$\ulcorner (\mathcal{B} *_\delta \alpha) *_\varepsilon \beta \urcorner = \ulcorner (\leq^*_{\alpha,\delta})^*_{\beta,\varepsilon} \urcorner$$
$$\ulcorner ((\mathcal{B} *_\delta \alpha) *_\varepsilon \beta) *_\zeta \gamma \urcorner = \ulcorner ((\leq^*_{\alpha,\delta})^*_{\beta,\varepsilon})^*_{\gamma,\zeta} \urcorner$$

... and so on.

The proof of this observation has to be supplied elsewhere.

3.4. Bounded revision as an operation on prioritized data bases

Space does not permit us to explain how the shifting of priorities in a prioritized data base $\overrightarrow{h} = h_1 \prec \ldots \prec h_n$ works. For this, the reader is referred to Rott (forthcoming). We reproduce the transition instruction here just in order to convey an idea of the complexity involved.

(BoundRevPDB)

$$\boxed{\overrightarrow{h} \quad \longmapsto \quad \overrightarrow{h} \prec. \; \alpha \prec. \; \overrightarrow{h_{\leq(\alpha\to\delta)} \vee \alpha} \prec. \; \overrightarrow{h_{>(\alpha\to\delta)}}}$$

4. Revision by comparison

'Revision by comparison' is the name of the second existing kind of two-dimensional belief change that we consider in this paper. Fermé and Rott (2004) present a gentle introduction into the intuitive ideas and formal properties of revision by comparison. Here is their set of axioms for the case of one-step (i.e., non-iterated) revision.

(C1) $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ is logically closed

(C2) If α is logically equivalent with β and δ is logically equivalent with ε, then $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\varepsilon \beta$

(C3) If ε is not in $\ulcorner \mathcal{B} *_\delta \bot \urcorner$, then $\ulcorner \mathcal{B} *_\delta \bot \urcorner$ is a subset of $\ulcorner \mathcal{B} *_\varepsilon \bot \urcorner$

(C4) If δ is in $\ulcorner \mathcal{B} *_\delta \bot \urcorner$, then δ is in every $\ulcorner \mathcal{B} *_\varepsilon \alpha \urcorner$

(C5) If δ is in $\ulcorner \mathcal{B} *_{\delta\wedge\neg\alpha} \bot \urcorner$, then $\ulcorner \mathcal{B} *_\delta \alpha \urcorner = Cn(\ulcorner \mathcal{B} *_{\delta\wedge\neg\alpha} \bot \urcorner + \alpha)$

(C5) If δ is not in $\ulcorner \mathcal{B} *_{\delta \wedge \neg\alpha} \bot \urcorner$, then $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner = \ulcorner \mathcal{B} *_{\delta \wedge \neg\alpha} \bot \urcorner$

Revision by comparison does not satisfy all the AGM axioms, but a certain variation of that set. The following conditions are theorems following from (C1)–(C6).[23]

(AGM1) $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$ is logically closed.

(AGM^{c}2) $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$ contains α, if it contains δ.

(AGM3) $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$ is a subset of $Cn(\ulcorner \mathcal{B} \urcorner + \alpha)$

(AGM4) If α is consistent with $\ulcorner \mathcal{B} \urcorner$, then $\ulcorner \mathcal{B} \urcorner$ is a subset of $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$

(AGM^{c}5) $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$ is inconsistent iff $\neg\alpha, \delta \in \ulcorner \mathcal{B} *_{\varepsilon} \beta \urcorner$ for all ε and β

(AGM6) If α is logically equivalent with β, then $\mathcal{B} *_{\delta} \alpha \simeq \mathcal{B} *_{\delta} \beta$

(AGM7) $\ulcorner \mathcal{B} *_{\delta} (\alpha \wedge \beta) \urcorner$ is a subset of $Cn(\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner + \beta)$

(AGM8) If β is consistent with $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$, then $\ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$ is a subset of $\ulcorner \mathcal{B} *_{\delta} (\alpha \wedge \beta) \urcorner$

(AGM7&8) $\mathcal{B} *_{\delta} (\alpha \vee \beta) \simeq \mathcal{B} *_{\delta} \alpha$ or $\mathcal{B} *_{\delta} (\alpha \vee \beta) \simeq \mathcal{B} *_{\delta} \beta$ or $\ulcorner \mathcal{B} *_{\delta} (\alpha \vee \beta) \urcorner = \ulcorner \mathcal{B} *_{\delta} \alpha \urcorner \cap \ulcorner \mathcal{B} *_{\delta} \beta \urcorner$

The restricted success condition (AGM^{c}2) is a lot weaker than the unconditional AGM success condition (AGM2). If the negation of the input sentence is at least as entrenched as the reference sentence, then the latter gets lost rather than the input accepted. (AGM^{c}5) is different from the corresponding AGM postulate in that it is an 'iff' rather than an 'if' condition, and it allows that a consistent input may lead to an inconsistent belief set, but only if both the input's negation and the reference sentence are irrevocable. Two further very strong conditions for varying reference sentences are satisfied in revision by comparison (Fermé and Rott 2004, Obs. 6):

(AGM$-^c$8) If $\delta \notin \ulcorner \mathcal{B} *_{\delta \wedge \varepsilon} \alpha \urcorner$, then $\ulcorner \mathcal{B} *_{\delta \wedge \varepsilon} \alpha \urcorner = \ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$

(AGM$-^c$D) $\ulcorner \mathcal{B} *_{\delta \wedge \varepsilon} \alpha \urcorner = \ulcorner \mathcal{B} *_{\delta} \alpha \urcorner$ or $\ulcorner \mathcal{B} *_{\delta \wedge \varepsilon} \alpha \urcorner = \ulcorner \mathcal{B} *_{\varepsilon} \alpha \urcorner$

Now we ask whether revision by comparison satisfies the Same beliefs condition (SBC). It is clear from the discussion of Fermé and Rott (2004) that it does not. If the reference sentence δ is strong enough, i.e., if it is more entrenched than the negation $\neg\alpha$ of the input sentence α, then revision by comparison yields an AGM revision of the initial belief set. If, however, the reference sentence is

[23] In (AGM3) and (AGM4), as well as in Observation 3 below, conditions (i) and (ii), $\mathcal{B}$ is identified with $\mathcal{B} *_{\bot} \bot$. For an explanation, see Rott and Fermé (2004, p. 19).

weaker, than it gets lost, and what we get is not a successful revision by α but a severe withdrawal (Pagnucco and Rott 1999) of the reference sentence δ.

Is it possible to characterize special cases in which (SBC) holds for revision by comparison? Although the building blocks were present in Fermé and Rott (2004), they did not address the problem head-on. So we supply the following

Observation 3. Revision by comparison satisfies the following conditions.

(i) If α is in $\ulcorner\mathcal{B}\urcorner$ and $\ulcorner\mathcal{B}\urcorner$ is consistent, then $\mathcal{B} *_\delta \alpha \simeq \mathcal{B}$.

(ii) If δ is not in $\ulcorner\mathcal{B}\urcorner$, then $\mathcal{B} *_\delta \alpha \simeq \mathcal{B}$.

(iii) $\mathcal{B} *_\delta \top \simeq \mathcal{B}$.

(iv) If α is in $\ulcorner\mathcal{B} *_\delta \alpha\urcorner$ and $\ulcorner\mathcal{B}\urcorner$ is consistent, then $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\top \alpha$

(v) If α is in both $\ulcorner\mathcal{B} *_\delta \alpha\urcorner$ and $\ulcorner\mathcal{B} *_\varepsilon \alpha\urcorner$, then $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\varepsilon \alpha$ (SBC^{c}1)

(vi) If α is not in $\ulcorner\mathcal{B} *_\delta \alpha\urcorner$, then $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\delta \bot$

(vii) If α is neither in $\ulcorner\mathcal{B} *_\delta \alpha\urcorner$ nor in $\ulcorner\mathcal{B} *_\varepsilon \alpha\urcorner$, then $\ulcorner\mathcal{B} *_\delta \alpha\urcorner$ is a subset or a superset of $\ulcorner\mathcal{B} *_\varepsilon \alpha\urcorner$

(viii) If neither δ nor ε is in $\ulcorner\mathcal{B} *_{\delta\wedge\varepsilon} \bot\urcorner$, then $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\varepsilon \alpha$ (SBC^{c}2)

The proof of Observation 3 is to be found in the appendix. Here are some comments. (i) is a vacuity condition for the input sentence, (ii) a vacuity condition for the reference sentence.[24] (iii) says that revising by a tautology does not change the belief set. (iv) tells us that as far as the belief sets are concerned, any successful revision by α gives the same result as an 'irrevocable revision' by α. (v) is a Same Belief Condition *restricted to successful revisions by* α; we call it (SBC^{c}1). (vi) tells us that as far as the belief sets are concerned, any unsuccessful revision by α give us the same result as a 'severe withdrawal' of the reference sentence δ. (vii) states that all belief sets reached by unsuccessful revisions by α are related by subset inclusion. (viii) is a Same Belief Condition *restricted to revisions with reference sentences of equal entrenchment*. This restriction is expressed by the antecedent of condition (viii) which we also call (SBC^{c}2).

4.1. Revision by comparison as an operation for iterated belief change

The general condition for the two-dimensional operation of *revision by comparison* is quite complicated (Fermé and Rott 2004):

[24] Condition (ii), but not condition (i) could even be strengthened to an identity of belief states rather than only of the belief sets.

(IT^c)

$$\ulcorner(\mathcal{B} *_\delta \alpha) *_\varepsilon \bot\urcorner \;=\; \begin{cases} Cn(\ulcorner\mathcal{B} \circ_{\alpha\to\varepsilon} \bot\urcorner + \alpha) & \text{if } \ulcorner\mathcal{B} *_\varepsilon \bot\urcorner \neq \mathcal{L} \\ & \quad \text{and } \;\delta \in \ulcorner\mathcal{B} \circ_\delta (\alpha \wedge \neg\varepsilon)\urcorner \\ \ulcorner\mathcal{B} *_\delta \bot\urcorner \;\cap\; \ulcorner\mathcal{B} *_\varepsilon \bot\urcorner & \text{if } \ulcorner\mathcal{B} *_\varepsilon \bot\urcorner \neq \mathcal{L} \\ & \quad \text{and } \;\delta \notin \ulcorner\mathcal{B} *_\delta (\alpha \wedge \neg\varepsilon)\urcorner \\ \mathcal{L} & \text{if } \ulcorner\mathcal{B} *_\varepsilon \bot\urcorner = \mathcal{L} \end{cases}$$

There are a number of simple derived properties for iteration that are satisfied by revision by comparison.

(IT^{c}1) $(\mathcal{B} *_\delta \alpha) *_\delta \beta \;\simeq\; \mathcal{B} *_\delta (\alpha \wedge \beta)$

(IT^{c}2) $(\mathcal{B} *_\delta \alpha) *_\varepsilon \alpha \;\simeq \begin{cases} \mathcal{B} *_\delta \alpha & \text{, if } \;\varepsilon \notin \ulcorner\mathcal{B} *_\delta \bot\urcorner \text{ or } \ulcorner\mathcal{B} *_\delta \bot\urcorner = \mathcal{L} \\ \mathcal{B} *_\varepsilon \alpha & \text{, otherwise} \end{cases}$

(IT^{c}3) $(\mathcal{B} *_\delta \alpha) *_\varepsilon \alpha \;\simeq\; (\mathcal{B} *_\varepsilon \alpha) *_\delta \alpha$

(IT^{c}4) $\ulcorner(\mathcal{B} *_\delta \bot) *_\varepsilon \bot\urcorner = \begin{cases} \ulcorner\mathcal{B} *_\delta \bot\urcorner \cap \ulcorner\mathcal{B} *_\varepsilon \bot\urcorner & \text{, if } \;\ulcorner\mathcal{B} *_\delta \bot\urcorner, \ulcorner\mathcal{B} *_\varepsilon \bot\urcorner \neq \mathcal{L} \\ \mathcal{L} & \text{, otherwise} \end{cases}$

(IT^{c}5) $(\mathcal{B} *_\delta \alpha) *_\alpha \delta \;\simeq \begin{cases} \mathcal{B} *_\alpha \delta & \text{, if } \;\delta \notin \ulcorner\mathcal{B} *_\alpha \bot\urcorner \text{ or } \ulcorner\mathcal{B} *_\alpha \bot\urcorner = \mathcal{L} \\ \mathcal{B} *_\delta \alpha & \text{, otherwise} \end{cases}$

Fermé and Rott (2004) provided an inelegant axiomatization, but no elegant axiomatization of iterated revision by comparison is known yet. In particular, it is not clear whether (IT^{c}1)–(IT^{c}5) jointly sufficient to derive (IT^c) and thus to characterize the iteration part of revision by comparison.

4.2. Revision by comparison as an operation on systems of spheres

Here the general questions are quite similar to those of Section 3.2. The particular s.o.s. semantics spelt out in Fermé and Rott is described by the following equation:

(RevCompSS)

$$\$^*_{\alpha,\delta} = \{S \cap [\alpha] : S \in \$, S \cap [\alpha] \neq \emptyset \text{ and } S \subseteq [\delta]\} \cup \{S : S \in \$ \text{ and } S \not\subseteq [\delta]\}$$

Figure 2 shows what happens to an s.o.s. when it gets revised by revision by comparison.

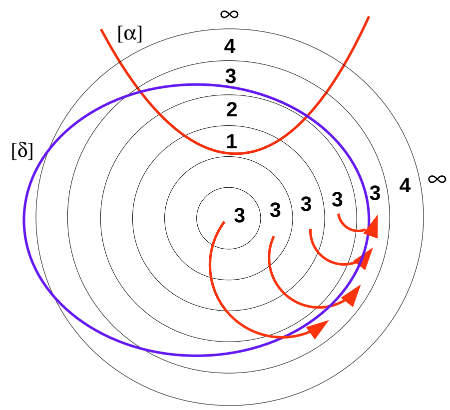

Fig. 2: Revision by comparison: Moving the best $\neg\alpha \wedge \delta$-worlds to the closest $\neg\delta$-permitting sphere, deleting all distinctions

The picture brings out the fact that revision by comparison tends to decrease the number of spheres in an s.o.s. (in the example from 6 to 4), thus making plausibility distinctions coarser. This contrasts with bounded revision that tends to increase the number of spheres and thus introduces finer plausibility distinctions. In a way, the two methods seem to complement each other.

A result analogous to Observation 3.2 holds true, but has not been proved anywhere so far. We refer the reader to the equivalent modelling in terms of entrenchment relations (see the next section).

4.3. Revision by comparison as an operation on entrenchment relations

The preliminaries for this section are similar to Section 3.3. The corresponding result for revision by comparison is this:

(RevCompEnt)

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

Results analogous to Observation 2 have been provided by Fermé and Rott (2004, Theorems 10–13).

4.4. Revision by comparison as an operation on prioritized data bases

As in the case of bounded revision, we refrain from explaining any details here (Rott forthcoming). We must be content with an impression of the complexity of the operation. Revision by comparison transforms a prioritized data base $\overrightarrow{h} = h_1 \prec \ldots \prec h_n$ into the following one:

(RevCompPDB)

$$\boxed{\overrightarrow{h} \quad \longmapsto \quad \overrightarrow{h_{<\delta}} \prec . \; h_{=\delta} \wedge \alpha \prec . \; \overrightarrow{h_{>\delta}}}$$

This is a surprisingly simple recipe. The input sentence α is inserted at the highest level in the prioritized data base at which the reference sentence is derivable.

5. Conclusion

We have discussed two two-dimensional operations of belief revision that 'lie between' quantitative and qualitative approaches in that they do not use numbers and are still able to specify the extent or degree to which a new piece of information is to be accepted. They do so by specifying a reference sentence with the idea that the input has to be accepted *as long as* the reference sentence holds along with the input sentence (and just a little further than that), or alternatively, *at least as* strongly as the reference sentence. In both senses, the acceptance of the input sentence may be said to be delimited by the range or the strength of the reference sentence.

The models compared are *bounded revision* as introduced in Rott (2007) and *revision by comparison* as studied by Cantwell (1997) and Fermé and Rott (2004).[25] Despite their principal similarities, these models exhibit a number of important differences, differences that make for a quite different *Gestalt*. We list eight features in roughly decreasing order of significance.

(1) Bounded revision is *successful* in the sense that the input sentence always gets accepted, independently of which reference sentence is used. Revision by comparison, in contrast, is successful only in the severely restricted form of (AGM^c2). If the reference sentence is not more entrenched than the negation

[25]Cantwell uses the term 'raising' and contrasts it with yet another two-dimensional operation he calls 'lowering'.

of the input sentence, then the former gets lost rather than the latter gets accepted.

(2) Bounded revision satisfies the *Same Beliefs Condition* (SBC) unconditionally, while revision by comparison satisfies it only provided that either the revision is 'successful' or the strength of the index sentence is invariant ((SBC^c1) and (SBC^c2)). The explanation for both (1) and (2) is, of course, that in contrast to bounded revision, revision by comparison embodies not only an operation of belief revision, but also an operation of belief contraction.[26]

(3) Both models fill out a whole space of possibilities between interesting one-dimensional belief change functions as limiting cases. But these limiting cases are quite different. Taking a logical truth as the reference sentence gives *irrevocable revision* (see Segerberg 1998 and Rott 2006) for revision by comparison, while it gives *moderate revision* for bounded revision. Taking a logical falsity as the reference sentence generates *conservative revision* for bounded revision, but does not produce any change for revision by comparison. Fixing a logical falsity as the input sentence produces a *severe withdrawal* of the reference sentence in revision by comparison, but does not produce any changes in the orderings representing the belief state.[27]

(4) Bounded revision satisfies the *Darwiche-Pearl postulates*. In contrast, revision by comparison violates these postulates, since it wipes out relevant distinctions between worlds in which the input sentence is false and the reference sentence is true.

(5) While bounded revision tends to *refine* orderings of possible worlds and beliefs (the number of spheres in the agent's s.o.s. and the number of layers in her entrenchment relation increase), revision by comparison has just the opposite effect and tends to *coarsen* orderings of possible worlds and beliefs (the number of spheres and entrenchment layers decreases).

(6) While bounded revision follows basically an *as-long-as strategy* ('accept α as long as δ holds along with it'), revision by comparison opts for an *at-least-as strategy* ('accept α as least as strongly as δ').

[26] In more precise terminology: an operation of severe withdrawal (Pagnucco and Rott 1999).

[27] In terms of systems of spheres, revising by an inconsistency according to (BoundRevSS) only adds the empty sphere as the new innermost sphere and thus generates an inconsistent belief set(but no change in the corresponding ordering of possible worlds). In terms of entrenchments, revising by an inconsistency according to (BoundRevEnt) introduces no change in the ordering of sentences.

(7) In bounded revision, the input sentence is believed in the posterior belief state *just a little more strongly* than the reference sentence, it is believed *at least as strongly* as the reference sentence in revision by comparison.

(8) Various presentations of the two operations give quite different views of their *complexity*. The iteration axiom (IT^b) for bounded revision is simple, but the iteration axiom (IT^c) for revision by comparison looks terribly complicated.[28] On the other hand, the presentation of revision by comparison in terms of changes of prioritized data bases is very simple, while the one of bounded revision is fairly complex.

There is a whole list of questions concerning further methods of two-dimensional belief change. Is it possible to use bounded revision with an at-least-as strategy?[29] Can we use it with a posterior of '$\delta \leq \alpha$' in the place of '$\delta < \alpha$'? Can we equip bounded revision with an integrated contraction mechanism? Or, what may amount to the same thing, can we retain the idea of revision by comparison and at the same time satisfy the Darwiche-Pearl postulates? Is it possible to combine revision by comparison with an as-long-as strategy? Question such as these indicate the richness contained in the idea of two-dimensional belief revision. Last but not least, we should not forget to mention that in the very first paper (I think) presenting the idea of two-dimensional belief change, John Cantwell (1997) introduced a method of 'lowering' that we have not covered at all in this paper. I am sure that many interesting discoveries can be made about a great diversity of two-dimensional belief change operations.

The eight factors listed above are certainly not independent of each other. But taken together they show that bounded revision and revision by comparison are complements in various interesting ways. Bounded revision and revision by comparison are two implementations of a single very general idea, that of renouncing the use of numbers and working with reference sentences in their place, thus interpreting belief change as a sort of doxastic preference change with inputs of the form '$\delta \leq \alpha$' or '$\delta < \alpha$'. There are some good reasons why just these two operations were chosen as objects of study, but they are definitely not the only reasonable ways to go two-dimensional. Going two-dimensional gives a lot of new leeway for approaches that refrain from assuming meaningful numbers as measuring degrees of belief. I hope to have indicated that one can work without numbers and advance to more elaborate forms of reasoning than

[28]As mentioned above, it is possible that (IT^c) can be replaced by one or more simple iteration axioms.

[29]The answer to this question is 'yes.'

the one reported by Gordon (2004).[30] This paper is meant as an invitation to cooperate and explore a rich diversity of possibilities in two-dimensional belief change. As research in belief revision progresses, an increasing number of potentially rational methods for revising one's belief states emerges. What we will need in order to apply these promising models in practice is a general methodology telling us when to apply which operations of belief change.

Appendix: Proof of Observation 3

Most of the properties listed in Observation 3 are immediate from the semantics of revision by comparison as given in Fermé and Rott. The proofs from the first principles (C1)–(C6) are more tedious, but still they are instructive. We will freely use the properties (Q1)–(Q16) from Fermé and Rott (2004, Obs. 7).

(i) Let α be in $\ulcorner\mathcal{B}\urcorner$. Then, by (AGM3), $\ulcorner\mathcal{B} *_\delta \alpha\urcorner \subseteq Cn(\ulcorner\mathcal{B}\urcorner + \alpha) = \ulcorner\mathcal{B}\urcorner$. Since $\alpha \in \ulcorner\mathcal{B}\urcorner \neq \mathcal{L}$, α is consistent with $\ulcorner\mathcal{B}\urcorner$. So, by (AGM4), $\ulcorner\mathcal{B}\urcorner \subseteq \ulcorner\mathcal{B} *_\delta \alpha\urcorner$. In sum thus $\mathcal{B} *_\delta \alpha \simeq \mathcal{B}$.

(ii) is Lemma 5 of Fermé and Rott (2004).

(iii) By (C5) and (C6), $\ulcorner\mathcal{B} *_\delta \top\urcorner$ is either $Cn(\ulcorner\mathcal{B} *_{\delta\wedge\neg\top} \bot\urcorner + \top)$ or $\ulcorner\mathcal{B} *_{\delta\wedge\neg\top} \bot\urcorner$, which both reduces to $\ulcorner\mathcal{B} *_\bot \bot\urcorner$. Since $\mathcal{B}$ is identified with $\mathcal{B} *_\bot \bot$, we get that $\mathcal{B} \simeq \mathcal{B} *_\bot \bot \simeq \mathcal{B} *_\delta \top$.

(iv) Let $\alpha \in \ulcorner\mathcal{B} *_\delta \alpha\urcorner$. Then, by (Q11), either $\delta \in \ulcorner\mathcal{B} *_\delta \alpha\urcorner$ or $\alpha \in \ulcorner\mathcal{B} *_\delta \neg\alpha\urcorner$. If the former is true, we conclude with (Q10) directly that $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\top \alpha$, and we are done. So suppose the latter is true. From this, together with $\alpha \in \ulcorner\mathcal{B} *_\delta \alpha\urcorner$, we can derive with the help of (AGM7&8) that $\alpha \in \ulcorner\mathcal{B} *_\delta (\alpha \vee \neg\alpha)\urcorner$ from which we get, by (C2), $\alpha \in \ulcorner\mathcal{B} *_\delta \top\urcorner$. Thus, by (iii), $\alpha \in \ulcorner\mathcal{B}\urcorner$, so since $\ulcorner\mathcal{B}\urcorner$ is assumed to be consistent, we can use (i) and conclude that $\mathcal{B} *_\delta \alpha \simeq \ulcorner\mathcal{B}\urcorner \simeq \mathcal{B} *_\top \alpha$, as desired.

(v) follows from (iv).

(vi) Suppose that $\alpha \notin \ulcorner\mathcal{B} *_\delta \alpha\urcorner$. Then by axioms (C5) and (C6) $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_{\delta\wedge\neg\alpha} \bot$. Also by (AGM^{c}2) $\delta \notin \ulcorner\mathcal{B} *_\delta \alpha\urcorner$. So $\delta \notin \ulcorner\mathcal{B} *_{\delta\wedge\neg\alpha} \bot\urcorner$ and by axiom (C3) $\ulcorner\mathcal{B} *_{\delta\wedge\neg\alpha} \bot\urcorner \subseteq \ulcorner\mathcal{B} *_\delta \bot\urcorner$. On the other hand, from $\delta \notin \ulcorner\mathcal{B} *_\delta \alpha\urcorner$ we also get, by axiom (C4), $\delta \notin \ulcorner\mathcal{B} *_\delta \bot\urcorner$. From this, property (Q16) gives $\ulcorner\mathcal{B} *_\delta \bot\urcorner \subseteq$

[30] Other cognitive activities without numerals and numbers are discussed by Field (1980) and Hellman (1989).

$\ulcorner \mathcal{B} *_{\delta \wedge \neg \alpha} \perp \urcorner$. Taking the two inclusions together, we get $\mathcal{B} *_\delta \perp \simeq \mathcal{B} *_{\delta \wedge \neg \alpha} \perp$. But we already know that $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_{\delta \wedge \neg \alpha} \perp$, and finally arrive at $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\delta \perp$, as desired.

(vii) Suppose that α is neither in $\ulcorner \mathcal{B} *_\delta \alpha \urcorner$ nor in $\ulcorner \mathcal{B} *_\varepsilon \alpha \urcorner$. Then we know from (vi) that $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\delta \perp$ and $\mathcal{B} *_\varepsilon \alpha \simeq \mathcal{B} *_\varepsilon \perp$. We distinguish two cases. First, suppose that $\delta \notin \ulcorner \mathcal{B} *_\varepsilon \perp \urcorner$. Then, by (C3), $\ulcorner \mathcal{B} *_\varepsilon \perp \urcorner \subset \ulcorner \mathcal{B} *_\delta \perp \urcorner$, and we are done. So suppose, second, that $\delta \in \ulcorner \mathcal{B} *_\varepsilon \perp \urcorner$. If also $\delta \in \ulcorner \mathcal{B} *_\delta \perp \urcorner$, then $\ulcorner \mathcal{B} *_\delta \perp \urcorner = \mathcal{L}$, by Lemma 0(d) of Fermé and Rott (2004), and $\ulcorner \mathcal{B} *_\varepsilon \perp \urcorner \subseteq \ulcorner \mathcal{B} *_\delta \perp \urcorner$ again. If, on the other hand, $\delta \notin \ulcorner \mathcal{B} *_\delta \perp \urcorner$, then, using (Q1), we immediately get $\ulcorner \mathcal{B} *_\delta \perp \urcorner \subseteq \ulcorner \mathcal{B} *_\varepsilon \perp \urcorner$, and we are done as well.

(viii) As a preparatory lemma, we show that for all δ and ε,

(+) Either $\ulcorner \mathcal{B} *_\delta \perp \urcorner \subseteq \ulcorner \mathcal{B} *_\varepsilon \perp \urcorner$ or $\ulcorner \mathcal{B} *_\varepsilon \perp \urcorner \subseteq \ulcorner \mathcal{B} *_\delta \perp \urcorner$.

(+) is immediate if either $\ulcorner \mathcal{B} *_\delta \perp \urcorner$ or $\ulcorner \mathcal{B} *_\varepsilon \perp \urcorner$ in inconsistent. So suppose that neither is inconsistent. By (Q16), we then have both $\ulcorner \mathcal{B} *_\delta \perp \urcorner \subseteq \ulcorner \mathcal{B} *_{\delta \wedge \varepsilon} \perp \urcorner$ and $\ulcorner \mathcal{B} *_\varepsilon \perp \urcorner \subseteq \ulcorner \mathcal{B} *_{\delta \wedge \varepsilon} \perp \urcorner$. From this, (+) follows immediately with the help of (AGM$-^c$D).

Suppose as the general hypothesis that neither δ nor ε is in $\ulcorner \mathcal{B} *_{\delta \wedge \varepsilon} \perp \urcorner$. Then, by (AGM$-^c$8), we get $\mathcal{B} *_\delta \perp \simeq \mathcal{B} *_{\delta \wedge \varepsilon} \perp \simeq \mathcal{B} *_\varepsilon \perp$. Suppose first that $\neg\alpha \in \ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner$. Then we prove that the claim $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\varepsilon \alpha$ is equivalent with $\mathcal{B} *_\delta \perp \simeq \mathcal{B} *_\varepsilon \perp$ which we have just shown. Let $\neg\alpha \in \ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner$. By (C4) then, $\neg\alpha \in \ulcorner \mathcal{B} *_\delta \alpha \urcorner \neq \mathcal{L}$, so $\alpha \notin \ulcorner \mathcal{B} *_\delta \alpha \urcorner$ and by (AGM8) $\ulcorner \mathcal{B} *_\delta \alpha \urcorner \subseteq \ulcorner \mathcal{B} *_\delta \perp \urcorner$. On the other hand, we have $\delta \notin \ulcorner \mathcal{B} *_\delta \perp \urcorner$, so by (Q17), $\ulcorner \mathcal{B} *_\delta \perp \urcorner \subseteq \ulcorner \mathcal{B} *_\delta \alpha \urcorner$. Taken together, we get $\ulcorner \mathcal{B} *_\delta \alpha \urcorner = \ulcorner \mathcal{B} *_\delta \perp \urcorner$, which is what we needed.

So suppose second that $\neg\alpha \notin \ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner$. We address (viii) by a case distinction that is complete, by (+).

Case 1: $\ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner \subseteq \ulcorner \mathcal{B} *_\delta \perp \urcorner$. By (AGM$-^c$D) and (Q16), we get $\mathcal{B} *_{\delta \wedge \neg\alpha} \perp \simeq \mathcal{B} *_\delta \perp$. Since $\delta \notin \ulcorner \mathcal{B} *_\delta \perp \urcorner$, we get from (C6) that $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_{\delta \wedge \neg\alpha} \perp \simeq \mathcal{B} *_\delta \perp$. Since we know from the hypothesis that $\mathcal{B} *_\varepsilon \perp \simeq \mathcal{B} *_\delta \perp$, we get from the assumption of Case 1 that $\ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner \subseteq \ulcorner \mathcal{B} *_\varepsilon \perp \urcorner$. Then by reasoning in a completely analogous manner as before, we get $\mathcal{B} *_\varepsilon \alpha \simeq \mathcal{B} *_{\varepsilon \wedge \neg\alpha} \perp \simeq \mathcal{B} *_\varepsilon \perp$. Using the transitivity of $\simeq$ to wrap things up, we finally arrive at $\mathcal{B} *_\delta \alpha \simeq \mathcal{B} *_\varepsilon \alpha$, as desired.

Case 2: $\ulcorner \mathcal{B} *_\delta \perp \urcorner \subset \ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner$. From this it follows that $\ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner \not\subseteq \ulcorner \mathcal{B} *_\delta \perp \urcorner$, and so by (C3), $\delta \in \ulcorner \mathcal{B} *_{\neg\alpha} \perp \urcorner$. On the other hand, from the assumption that

$\ulcorner\mathcal{B}*_{\delta}\bot\urcorner \subseteq \ulcorner\mathcal{B}*_{\neg\alpha}\bot\urcorner$, we get by (AGM$-^{c}$D) and (Q16) that $\mathcal{B}*_{\delta\wedge\neg\alpha}\bot \simeq \mathcal{B}*_{\neg\alpha}\bot$. Thus $\delta \in \ulcorner\mathcal{B}*_{\delta\wedge\neg\alpha}\bot\urcorner$, and so by (C5), we get that $\ulcorner\mathcal{B}*_{\delta}\alpha\urcorner = Cn(\ulcorner\mathcal{B}*_{\delta\wedge\neg\alpha}\bot\urcorner + \alpha) = Cn(\ulcorner\mathcal{B}*_{\neg\alpha}\bot\urcorner + \alpha) = \ulcorner\mathcal{B}*_{\top}\alpha\urcorner$, by (C5) again. Since we know from the hypothesis that $\mathcal{B}*_{\varepsilon}\bot \simeq \mathcal{B}*_{\delta}\bot$, we get from the assumption of Case 2 that $\ulcorner\mathcal{B}*_{\varepsilon}\bot\urcorner \subset \ulcorner\mathcal{B}*_{\neg\alpha}\bot\urcorner$. Then by reasoning in a completely analogous manner as before, we get $\ulcorner\mathcal{B}*_{\varepsilon}\alpha\urcorner = Cn(\ulcorner\mathcal{B}*_{\varepsilon\wedge\neg\alpha}\bot\urcorner + \alpha) = Cn(\ulcorner\mathcal{B}*_{\neg\alpha}\bot\urcorner + \alpha) = \ulcorner\mathcal{B}*_{\top}\alpha\urcorner$. In sum, we get again $\mathcal{B}*_{\delta}\alpha \simeq \mathcal{B}*_{\varepsilon}\alpha$, as desired. This completes the proof of (viii). QED

References

Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

Craig Boutilier. Revision sequences and nested conditionals. In Ruzena Bajcsy, editor, *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'93)*, pages 519–525, 1993.

John Cantwell. On the logic of small changes in hypertheories. *Theoria*, 63:54–89, 1997.

Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

Eduardo Fermé and Hans Rott. Revision by comparison. *Artificial Intelligence*, 157:5–47, 2004.

Hartry Field. *Science without numbers: A defence of nominalism.* Princeton University Press/Blackwell, Princeton, N.J./Oxford, 1980.

Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In Moshe Vardi, editor, *TARK'88: Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 83–95. Morgan Kaufmann, Los Altos, CA, 1988.

Peter Gordon. Numerical cognition without words: Evidence from Amazonia. *Science*, 306 (issue 5695):496–499, 2004.

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

Geoffrey Hellman. *Mathematics without numbers: Towards a modal-structural interpretation*. Clarendon Press, Oxford, 1989.

Abhaya Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41:353–390, 1994.

Maurice Pagnucco and Hans Rott. Severe withdrawal – and recovery. *Journal of Philosophical Logic*, 28:501–547, 1999. (correct, complete reprint with original pagination in the February 2000 issue; see publisher's 'Erratum', Journal of Philosophical Logic 29, 2000, p. 121.).

Hans Rott. *Change, Choice and Inference: A Study in Belief Revision and Nonmonotonic Reasoning*. Oxford University Press, Oxford, 2001.

Hans Rott. Basic entrenchment. *Studia Logica*, 73:257–280, 2003.

Hans Rott. Coherence and conservatism in the dynamics of belief. Part II: Iterated belief change without dispositional coherence. *Journal of Logic and Computation*, 13:111–145, 2003.

Hans Rott. Revision by comparison as a unifying framework: Severe withdrawal, irrevocable revision and irrefutable revision. *Theoretical Computer Science*, 355:228–242, 2006.

Hans Rott. Bounded revision: Two-dimensional belief change between conservatism and moderation. In Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson, and Dan Egonsson, editors, *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. Lunds Universitet, 2007. Available at www.fil.lu.se/hommageawlodek.

Hans Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In David Makinson, Jacek Malinowski, and Heinrich Wansing, editors, *Towards Mathematical Philosophy*, Trends in Logic. Springer Verlag, Berlin, forthcoming.

Krister Segerberg. Irrevocable belief revision in dynamic doxastic logic. *Notre Dame Journal of Formal Logic*, 39:287–306, 1998.