

Belief Revision I: The AGM Theory

Franz Huber

Department of Philosophy
University of Toronto, Canada

franz.huber@utoronto.ca

<http://huber.blogs.chass.utoronto.ca/>

penultimate version: please cite the paper in the
Philosophy Compass

July 7, 2014

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 3 |
| 2 | The AGM Theory of Belief Revision | 3 |
| 3 | Iterated Belief Revision | 10 |

Abstract

Belief revision theory studies how an ideal epistemic agent should revise her beliefs when she receives new information. In part I I will first present the AGM theory of belief revision (Alchourrón & Gärdenfors & Makinson 1985). Then I will focus on the problem of *iterated* belief revisions.

1 Introduction

Belief revision theory studies how an ideal epistemic agent should revise her beliefs when she receives new information. Sophia believes many things, among others that it will rain on Tuesday, that it will be sunny on Wednesday, and that weather forecasts are always reliable. Belief revision theory tells Sophia how to revise her beliefs when she learns that the weather forecast for Tuesday and Wednesday predicts rain. As we will see, this depends on the details of her beliefs, but under one way of filling in the details she should keep her belief that it will rain on Tuesday and give up her belief that it will be sunny on Wednesday. To state in full detail how Sophia should revise her beliefs when she learns new information we need a representation of her old beliefs and of the new information she receives.

Belief revision theory models belief as a qualitative attitude towards sentences or propositions: an ideal epistemic agent believes a proposition, or she disbelieves the proposition by believing its negation, or she suspends judgment with respect to the proposition and its negation. This is different in the theory of subjective probabilities, which is sometimes referred to as *Bayesianism* (Easwaran 2011a, 2011b). Here belief is modeled as a quantitative attitude towards a sentence or proposition: an ideal epistemic agent believes a proposition to a specific degree, her degree of belief or credence for the proposition. As we will see, though, in order to adequately model conditional beliefs and iterated belief revisions, belief revision theory has to model the ideal epistemic agent's beliefs with more than just the set of sentences or propositions she believes.

2 The AGM Theory of Belief Revision

The AGM theory of belief revision derives its name from the seminal paper by Alchourrón & Gärdenfors & Makinson (1985). Comprehensive overviews can be found in Gärdenfors (1988), Gärdenfors & Rott (1995), and Rott (2001).

One version of the AGM theory of belief revision represents the agent's old beliefs with a set of sentences from some formal language, her *belief set*, together with an *entrenchment ordering* over these sentences that represents how strongly the agent holds the beliefs in her belief set. The entrenchment ordering represents the details of the agent's beliefs. The new information is represented by a single sentence. The theory distinguishes between the easy case, called *expansion*, where the new information does not contradict the agent's old belief set and is

simply added conjunctively and the general case, called *revision*, where the new information might contradict the old belief set. The general case is difficult, because the agent has to turn her old belief set, which is assumed to be consistent, into a new belief set that contains the new information and is consistent. Usually the general case is dealt with in two steps: in a first step, called *contraction*, the old belief set is cleared from everything that contradicts the new information. In a second step one simply expands by adding the new information conjunctively. So the hard work is done by contraction which turns the general case of revision into the easy case of expansion.

A formal language \mathcal{L} is defined inductively as follows. \mathcal{L} contains the contradictory sentence \perp and all elements of a given countable set of propositional variables $PV = \{p, q, r, \dots\}$. Furthermore, whenever α and β are sentences of \mathcal{L} , then so are the negation of α , $\neg\alpha$, as well as the conjunction of α and β , $\alpha \wedge \beta$. Finally, nothing else is a sentence of \mathcal{L} .

The new information is represented by a single sentence α from \mathcal{L} . The ideal epistemic agent's old beliefs are represented by a set of sentences, her belief set $\mathcal{B} \subseteq \mathcal{L}$, and an entrenchment ordering \leq for \mathcal{B} . Where \vdash is the consequence relation on \mathcal{L} and $Cn(\mathcal{B}) = \{\alpha \in \mathcal{L} : \mathcal{B} \vdash \alpha\}$ the set of consequences of \mathcal{B} , the entrenchment ordering has to satisfy the following postulates. For all $\alpha, \beta, \gamma \in \mathcal{L}$:

- ≤ 1 If $\alpha \leq \beta$ and $\beta \leq \gamma$, then $\alpha \leq \gamma$. Transitivity
- ≤ 2 If $\beta \in Cn(\{\alpha\})$, then $\alpha \leq \beta$. Dominance
- ≤ 3 $\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$. Conjunctivity
- ≤ 4 If $\perp \notin Cn(\mathcal{B})$, then: $\alpha \notin \mathcal{B}$ just in case for all $\beta \in \mathcal{L}$: $\alpha \leq \beta$. Minimality
- ≤ 5 If for all α from \mathcal{L} : $\alpha \leq \beta$, then $\beta \in Cn(\emptyset)$. Maximality

Expansion, revision, and contraction turn belief sets and new information into belief sets. Formally they are functions from $\wp(\mathcal{L}) \times \mathcal{L}$ into $\wp(\mathcal{L})$. Expansion $\dot{+}$ turns each old belief set $\mathcal{B} \subseteq \mathcal{L}$ and sentence $\alpha \in \mathcal{L}$ into the new belief set $\mathcal{B} \dot{+} \alpha = Cn(\mathcal{B} \cup \{\alpha\})$. Revision $*$ turns each old belief set $\mathcal{B} \subseteq \mathcal{L}$ and sentence $\alpha \in \mathcal{L}$ into a new belief set $\mathcal{B} * \alpha$, where $*$ is required to satisfy the following postulates. For all $\mathcal{B} \subseteq \mathcal{L}$ and all $\alpha, \beta \in \mathcal{L}$:

- *1. $\mathcal{B} * \alpha = Cn(\mathcal{B} * \alpha)$. Closure
- *2. $\alpha \in \mathcal{B} * \alpha$. Success

- *3. $\mathcal{B} * \alpha \subseteq Cn(\mathcal{B} \cup \{\alpha\})$. Inclusion
- *4. If $\neg\alpha \notin Cn(\mathcal{B})$, then $\mathcal{B} \subseteq \mathcal{B} * \alpha$. Preservation
- *5. If $Cn(\{\alpha\}) = Cn(\{\beta\})$, then $\mathcal{B} * \alpha = \mathcal{B} * \beta$. Congruence
- *6. If $\neg\alpha \notin Cn(\emptyset)$, then $\perp \notin \mathcal{B} * \alpha$. Consistency
- *7. $\mathcal{B} * (\alpha \wedge \beta) \subseteq Cn((\mathcal{B} * \alpha) \cup \{\beta\})$. Conjunction 1
- *8. If $\neg\beta \notin \mathcal{B} * \alpha$, then $Cn((\mathcal{B} * \alpha) \cup \{\beta\}) \subseteq \mathcal{B} * (\alpha \wedge \beta)$. Conjunction 2

The two-step view of revision according to which it consists in first contracting the old belief set by the negation of the new information, thus making it consistent with the new information (as well as everything implied by the new information), and then expanding it by adding the new information conjunctively is described by the *Levi identity* (Levi 1977):

$$\mathcal{B} * \alpha = Cn(\mathcal{B} \dot{-} \neg\alpha \cup \{\alpha\})$$

The Levi identity puts contraction $\dot{-}$ center stage of the revision process. Contraction $\dot{-}$ turns each old belief set $\mathcal{B} \subseteq \mathcal{L}$ and sentence $\alpha \in \mathcal{L}$ into a reduced belief set $\mathcal{B} \dot{-} \alpha$ that is cleared from α (as well as everything implying α) and required to satisfy the following postulates. For all $\mathcal{B} \subseteq \mathcal{L}$ and all $\alpha, \beta \in \mathcal{L}$:

- $\dot{-}$ 1. $\mathcal{B} \dot{-} \alpha = Cn(\mathcal{B} \dot{-} \alpha)$. Closure
- $\dot{-}$ 2. If $\alpha \notin Cn(\emptyset)$, then $\alpha \notin Cn(\mathcal{B} \dot{-} \alpha)$. Success
- $\dot{-}$ 3. $\mathcal{B} \dot{-} \alpha \subseteq Cn(\mathcal{B})$. Inclusion
- $\dot{-}$ 4. If $\alpha \notin Cn(\mathcal{B})$, then $\mathcal{B} \dot{-} \alpha = \mathcal{B}$. Vacuity
- $\dot{-}$ 5. If $Cn(\{\alpha\}) = Cn(\{\beta\})$, then $\mathcal{B} \dot{-} \alpha = \mathcal{B} \dot{-} \beta$. Congruence
- $\dot{-}$ 6. $Cn(\mathcal{B}) \subseteq Cn((\mathcal{B} \dot{-} \alpha) \cup \{\alpha\})$. Recovery
- $\dot{-}$ 7. $(\mathcal{B} \dot{-} \alpha) \cap (\mathcal{B} \dot{-} \beta) \subseteq \mathcal{B} \dot{-} (\alpha \wedge \beta)$. Conjunction 1
- $\dot{-}$ 8. If $\alpha \notin \mathcal{B} \dot{-} (\alpha \wedge \beta)$, then $\mathcal{B} \dot{-} (\alpha \wedge \beta) \subseteq \mathcal{B} \dot{-} \alpha$. Conjunction 2

The Levi identity turns each contraction operator $\dot{-}$ satisfying $\dot{-}1\text{-}\dot{-}8$ into a revision operator $*$ satisfying $*1\text{-}*8$. The converse is true of the *Harper identity* (Harper 1976):

$$\mathcal{B}\dot{-}\alpha = \mathcal{B} * \neg\alpha \cap Cn(\mathcal{B})$$

Given a fixed belief set \mathcal{B} we can use an entrenchment ordering \leq for \mathcal{B} to define a revision operator $*$ as follows:

$$\mathcal{B} * \alpha = Cn(\{\beta \in \mathcal{L} : \neg\alpha < \beta\} \cup \{\alpha\}),$$

where $\alpha < \beta$ holds just in case $\alpha \leq \beta$ and $\beta \not\leq \alpha$.

The idea behind this equation is the following: when the agent revises $*$ her old belief set \mathcal{B} by the new information α she first has to clear \mathcal{B} from $\neg\alpha$ as well as everything else that is as entrenched as, or less entrenched than, $\neg\alpha$. For instance, \mathcal{B} also has to be cleared from everything implying $\neg\alpha$. However, all sentences β from the agent's old belief set \mathcal{B} that are more entrenched than $\neg\alpha$ can be preserved. This gives us the preserved belief set $\{\beta \in \mathcal{L} : \neg\alpha < \beta\}$. Then one adds the new information α to obtain $\{\beta \in \mathcal{L} : \neg\alpha < \beta\} \cup \{\alpha\}$. Finally one adds everything that is implied by the preserved belief set together with the new information.

Then one can prove

Theorem 1 *Let \mathcal{L} be a formal language and let $\mathcal{B} \subseteq \mathcal{L}$ be a belief set. For each entrenchment ordering \leq for \mathcal{B} satisfying $\leq 1\text{-}\leq 5$ there is a revision operator $*$ satisfying $*1\text{-}*8$ such that $\mathcal{B} * \alpha = Cn(\{\beta \in \mathcal{L} : \neg\alpha < \beta\} \cup \{\alpha\})$. For each revision operator $*$ satisfying $*1\text{-}*8$ there is an entrenchment ordering \leq for \mathcal{B} satisfying $\leq 1\text{-}\leq 5$ such that $\mathcal{B} * \alpha = Cn(\{\beta \in \mathcal{L} : \neg\alpha < \beta\} \cup \{\alpha\})$.*

See Gärdenfors (1988) and Gärdenfors & Makinson (1988).

There is a different way of representing postulates $*1\text{-}*8$ for revision operators $*$ due to Grove (1988). Similar to Lewis' (1973) theory of counterfactuals it uses systems of spheres defined on a set of possible worlds instead of entrenchment orderings defined on a formal language. A set of possible worlds can be thought of as a set of complete or maximally specific ways the world might be. One approach, used by Grove (1988), is to identify possible worlds with maximally consistent sets of sentences from \mathcal{L} , i.e. sets of sentences that are consistent, but that become inconsistent as soon as a single new sentence is added. Another approach is to take possible worlds as primitive. For present purposes we do not have to take a stance on this issue and can assume that we are given a set of possible worlds $W_{\mathcal{L}}$ relative to which we interpret the sentences from \mathcal{L} .

$\llbracket \alpha \rrbracket = \{w \in W_{\mathcal{L}} : w \models \alpha\}$ is the proposition expressed by the sentence α , i.e. the set of possible worlds in which α is true. If we follow Grove (1988) and construct $W_{\mathcal{L}}$ as the set of maximally consistent sets of sentences from \mathcal{L} , then $w \models \alpha$ holds just in case $\alpha \in w$. If possible worlds are taken as primitive, \models has to be analyzed differently. Either way there may be propositions $B \subseteq W_{\mathcal{L}}$ that are not expressed by any sentence α from \mathcal{L} . $\llbracket \mathcal{B} \rrbracket = \{w \in W_{\mathcal{L}} : w \models \alpha \text{ for all } \alpha \in \mathcal{B}\}$ is the proposition expressed by the set of sentences \mathcal{B} . If we follow Grove (1988) and construct $W_{\mathcal{L}}$ as the set of maximally consistent sets of sentences from \mathcal{L} , then for each proposition $B \subseteq W_{\mathcal{L}}$ there is a set of sentences from \mathcal{L} , $t(B)$, such that $B = \llbracket t(B) \rrbracket$ (it is $\bigcap \{\mathcal{A} \subseteq \mathcal{L} : \mathcal{A} \text{ is maximally consistent and } \mathcal{A} \in B\}$). This means that every proposition $B \subseteq W_{\mathcal{L}}$ can be expressed by a set of sentences from \mathcal{L} . If possible worlds are taken as primitive this has to be stipulated. We assume this stipulation to be made.

Let $B \subseteq W_{\mathcal{L}}$ be a proposition and let $\mathbf{S} \subseteq \wp(W_{\mathcal{L}})$ be a set of propositions. \mathbf{S} is a *system of spheres* in $W_{\mathcal{L}}$ that is *centered on* B if and only if for all $A, C \subseteq W_{\mathcal{L}}$ and all $\alpha \in \mathcal{L}$:

S1 If $A, C \in \mathbf{S}$, then $A \subseteq C$ or $C \subseteq A$. **S** is nested

S2 $B \in \mathbf{S}$, and: if $A \in \mathbf{S}$, then $B \subseteq A$. **S** is centered on B

S3 $W_{\mathcal{L}} \in \mathbf{S}$.

S4 If $\llbracket \alpha \rrbracket \cap D \neq \emptyset$ for some $D \in \mathbf{S}$, then there is $D^* \in \mathbf{S}$ such that: $\llbracket \alpha \rrbracket \cap D^* \neq \emptyset$, and $D^* \subseteq E$ for all $E \in \mathbf{S}$ with $\llbracket \alpha \rrbracket \cap E \neq \emptyset$.¹

In combination with **S3** requirement **S4** says that for each consistent sentence α there is a smallest sphere $D^* \in \mathbf{S}$ intersecting $\llbracket \alpha \rrbracket$. Let $c_{\mathbf{S}}(\alpha) := \llbracket \alpha \rrbracket \cap D^*$ and define $c_{\mathbf{S}}(\alpha) = \emptyset$ if α is inconsistent. $c_{\mathbf{S}}(\alpha)$ is the set of α -worlds that are “closest” to the center B , where closeness is spelt out in terms of the system of spheres \mathbf{S} . If α is consistent with (a set of sentences expressing) the center B , then $c_{\mathbf{S}}(\alpha)$ is just the intersection of the center B with the set of α -worlds, $\llbracket \alpha \rrbracket \cap B$. Otherwise one has to leave the center and move to the first sphere D^* that intersects with the proposition expressed by α and take their intersection, $\llbracket \alpha \rrbracket \cap D^*$, as $c_{\mathbf{S}}(\alpha)$.

¹In one respect Grove’s (1988) notion of a system of spheres is more general than Lewis’ (1973) notion as it allows \mathbf{S} to be centered on arbitrary propositions $B \subseteq W_{\mathcal{L}}$, whereas Lewis (1973: 14f) requires the center B to contain the actual world (and nothing but the actual world). In another respect Grove’s (1988) notion is less general than Lewis’ (1973) as requirement **S4** effectively makes the limit assumption, which Lewis (1973: 19f) famously rejects.

Given a fixed belief set \mathcal{B} we can use a system of spheres \mathcal{S} centered on $\llbracket \mathcal{B} \rrbracket$ to define a revision operator $*$ as follows:


$$\mathcal{B} * \alpha = t(c_{\mathcal{S}}(\alpha))$$

The idea is that what you end up believing after revising $*$ your old beliefs \mathcal{B} with the new information α is (a set of sentences expressing) the proposition $c_{\mathcal{S}}(\alpha)$ that contains the α -worlds that are closest when the proposition expressed by your old beliefs, $\llbracket \mathcal{B} \rrbracket$, is the center. Expansion is the special case where the proposition expressed by the new information intersects the proposition expressed by the old beliefs, $\llbracket \alpha \rrbracket \cap \llbracket \mathcal{B} \rrbracket \neq \emptyset$. In the general case of revision this intersection may be empty so that one has to leave the innermost sphere $\llbracket \mathcal{B} \rrbracket$ and move to the smallest sphere D^* that intersects with $\llbracket \alpha \rrbracket$.

Grove (1988) proves

Theorem 2 *Let \mathcal{L} be a formal language, let $\mathcal{B} \subseteq \mathcal{L}$ be a belief set, and let $W_{\mathcal{L}}$ be a set of possible worlds relative to which the sentences from \mathcal{L} are interpreted. For each system of spheres \mathcal{S} in $W_{\mathcal{L}}$ that is centered on $\llbracket \mathcal{B} \rrbracket$ and satisfies S1-S4 there is a revision operator $*$ satisfying *1-*8 such that $\mathcal{B} * \alpha = t(c_{\mathcal{S}}(\alpha))$. For each revision operator $*$ satisfying *1-*8 there is a system of spheres \mathcal{S} in $W_{\mathcal{L}}$ that is centered on $\llbracket \mathcal{B} \rrbracket$ and satisfies S1-S4 such that $\mathcal{B} * \alpha = t(c_{\mathcal{S}}(\alpha))$.*

We can picture a system of spheres centered on $\llbracket \mathcal{B} \rrbracket$ as an “onion” around $\llbracket \mathcal{B} \rrbracket$. The shaded area $\llbracket \mathcal{B} * \alpha \rrbracket$ is the strongest proposition the agent believes after revising her old belief set \mathcal{B} by the new information α :



In the AGM theory of belief revision the ideal epistemic agent's old beliefs are represented by her belief set \mathcal{B} together with her entrenchment ordering \leq for \mathcal{B} . The latter ordering guides the revision process in that it specifies which elements of the old belief set are given up and which are kept when new information δ is received. The result of revising the old belief set by new information is a new belief set $\mathcal{B} * \delta$. Sophia's old belief set \mathcal{B} includes the beliefs that it will rain on Tuesday, that it will be sunny on Wednesday, and that weather forecasts are always reliable. Suppose her belief α that it will be sunny on Wednesday is less entrenched than her belief β that it will rain on Tuesday, which in turn is less entrenched than her belief γ that weather forecasts are always reliable, $\alpha < \beta < \gamma$.

On Monday Sophia learns that the weather forecast for Tuesday and Wednesday predicts rain, δ . Consequently she has to give up her belief α that it will be sunny on Wednesday or her belief γ that weather forecasts are always reliable. The reason is that it follows from δ that one of those two beliefs is false, $\neg\alpha \vee \neg\gamma \in Cn(\{\delta\})$. This implies that $\alpha \wedge \gamma \leq \neg\delta$. Since $\beta \in Cn(\{\gamma, \delta\})$ there is no need for Sophia to give up her belief β that it will rain on Tuesday, and we assume $\neg\delta < \beta$ so that Sophia's entrenchment ordering looks as follows: $\alpha \sim \alpha \wedge \gamma \leq \neg\delta < \beta < \gamma$, where $x \sim y$ holds just in case $x \leq y$ and $y \leq x$. Thus Sophia's new belief set is $\mathcal{B} * \delta = Cn(\{x : \neg\delta < x\} \cup \{\delta\}) = Cn(\{\beta, \gamma, \delta, \neg\alpha\})$.

To Sophia's surprise it is sunny on Tuesday after all. Therefore Sophia wants to revise her newly acquired belief set $\mathcal{B} * \delta$ a second time by $\neg\beta$ to correct her belief β that it will rain on Tuesday. In addition, Sophia has to give up her belief δ that the weather forecast for Tuesday and Wednesday predicts rain (she might have misheard the weather forecast, say) or her belief γ that weather forecasts are always reliable (she might have been too gullible, say). The reason is that it follows from $\neg\beta$ that one of those two beliefs is false, $\neg\delta \vee \neg\gamma \in Cn(\neg\beta)$. Unfortunately AGM belief revision theory is of no help here. While Sophia could use her entrenchment ordering to revise her old belief set \mathcal{B} to a new belief set $\mathcal{B} * \delta$, the entrenchment ordering itself has not been updated and so is silent as to whether δ is more entrenched than γ (she was too gullible, say) or γ is more entrenched than δ (she misheard the weather forecast, say) or γ is as entrenched as δ (she was too gullible and misheard the weather forecast, say). This, however, is exactly the kind of information that Sophia would need to revise her beliefs a second time.

3 Iterated Belief Revision

More generally, the problem is that Sophia's beliefs are represented as a belief set plus an entrenchment ordering before the revision process, but as a belief set without an entrenchment ordering after the revision process. To handle *iterated* belief revisions the ideal epistemic agent's beliefs have to be represented in the same way before and after the revision process. Gärdenfors & Rott (1995: 37) call this the "principle of categorical matching".

Nayak (1994), Boutilier (1996), Darwiche & Pearl (1997), Segerberg (1998) and Fermé (2000), Rott (2003), Rott (2006), and others do exactly this. They augment the AGM postulates by further postulates specifying how the ideal epistemic agent should revise her entrenchment ordering in addition to her belief set when she receives new information. On their accounts the agent's beliefs are represented as a belief set plus an entrenchment ordering both before and after the revision process, and both of these two elements are revised when new information is received.

Let us have a closer look at the proposal by Darwiche & Pearl (1997). In addition to *1-*8 they propose the following four postulates for iterated belief revision. For all $\mathcal{B} \subseteq \mathcal{L}$ and all $\alpha, \beta \in \mathcal{L}$:

- *9. If $\beta \in Cn(\{\alpha\})$, then $(\mathcal{B} * \beta) * \alpha = \mathcal{B} * \alpha$.
- *10. If $\neg\beta \in Cn(\{\alpha\})$, then $(\mathcal{B} * \beta) * \alpha = \mathcal{B} * \alpha$.
- *11. If $\beta \in \mathcal{B} * \alpha$, then $\beta \in (\mathcal{B} * \beta) * \alpha$.
- *12. If $\neg\beta \notin \mathcal{B} * \alpha$, then $\neg\beta \notin (\mathcal{B} * \beta) * \alpha$.

In order to perspicuously represent these four postulates it is helpful to consider the following reformulation of a system of spheres \mathbf{S} in $W_{\mathcal{L}}$ centered on B . Let $B \subseteq W_{\mathcal{L}}$ be a proposition and let \leq be a binary relation on $W_{\mathcal{L}}$. \leq is a plausibility ordering on $W_{\mathcal{L}}$ with center B just in case for all $w, w', w'' \in W_{\mathcal{L}}$ and $A \subseteq W_{\mathcal{L}}$:

- ≤ 1 $w \leq w'$ or $w' \leq w$. \leq is connected
- ≤ 2 If $w \leq w'$ and $w' \leq w''$, then $w \leq w''$. \leq is transitive
- ≤ 3 $w \in B$ if and only if $\forall w^* : w \leq w^*$.
- ≤ 4 If $A \neq \emptyset$, then $\{v \in A : v \leq w^* \forall w^* \in A\} \neq \emptyset$.

So a system of spheres centered on B can be understood as a plausibility ordering with the center B being the set of most plausible possible worlds.

In terms of such a plausibility ordering the problem with the original AGM approach is the following. Before the revision process the agent's beliefs are represented as a belief set \mathcal{B} plus a plausibility ordering $\leq_{\mathcal{B}}$ with center $\llbracket \mathcal{B} \rrbracket$. After the revision process the agent's beliefs are represented as a belief set $\mathcal{B} * \alpha$, but without a corresponding plausibility ordering $\leq_{\mathcal{B} * \alpha}$. The principal of categorical matching urges us to represent the agent's beliefs as a belief set plus a plausibility ordering both before and after the revision process. In these terms postulates *9-*12 say the following:

- ≤ 5 . If $w, w' \in \llbracket \alpha \rrbracket$, then $w \leq_{\mathcal{B}} w'$ just in case $w \leq_{\mathcal{B} * \alpha} w'$.
- ≤ 6 . If $w, w' \notin \llbracket \alpha \rrbracket$, then $w \leq_{\mathcal{B}} w'$ just in case $w \leq_{\mathcal{B} * \alpha} w'$.
- ≤ 7 . If $w \in \llbracket \alpha \rrbracket$ and $w' \notin \llbracket \alpha \rrbracket$, and if $w <_{\mathcal{B}} w'$, then $w <_{\mathcal{B} * \alpha} w'$.
- ≤ 8 . If $w \in \llbracket \alpha \rrbracket$ and $w' \notin \llbracket \alpha \rrbracket$, and if $w \leq_{\mathcal{B}} w'$, then $w \leq_{\mathcal{B} * \alpha} w'$.

Here $w < w'$ holds just in case $w \leq w'$ and $w' \not\leq w$.

Grove (1988) proves

Theorem 3 *Let \mathcal{L} be a formal language, let $\mathcal{B} \subseteq \mathcal{L}$ be a belief set, and let $W_{\mathcal{L}}$ be a set of possible worlds relative to which the sentences from \mathcal{L} are interpreted. For each plausibility ordering \leq on $W_{\mathcal{L}}$ with center $\llbracket \mathcal{B} \rrbracket$ that satisfies ≤ 1 - ≤ 4 there is a revision operator $*$ satisfying *1-*8 such that $\mathcal{B} * \alpha = t(\{v \in \llbracket \alpha \rrbracket : v \leq w \forall w \in \llbracket \alpha \rrbracket\})$. For each revision operator $*$ that satisfies *1-*8 there is a plausibility ordering \leq on $W_{\mathcal{L}}$ with center $\llbracket \mathcal{B} \rrbracket$ satisfying ≤ 1 - ≤ 4 such that $\mathcal{B} * \alpha = t(\{v \in \llbracket \alpha \rrbracket : v \leq w \forall w \in \llbracket \alpha \rrbracket\})$.*

Darwiche & Pearl (1997) prove

Theorem 4 *Let \mathcal{L} be a formal language, let $\mathcal{B} \subseteq \mathcal{L}$ be a belief set, and let $W_{\mathcal{L}}$ be a set of possible worlds relative to which the sentences from \mathcal{L} are interpreted. Suppose $*$ satisfies *1-*8 and $\leq_{\mathcal{B}}$ is the corresponding plausibility ordering on $W_{\mathcal{L}}$ with center $\llbracket \mathcal{B} \rrbracket$ such that ≤ 1 - ≤ 4 and $\mathcal{B} * \alpha = t(\{v \in \llbracket \alpha \rrbracket : v \leq w \forall w \in \llbracket \alpha \rrbracket\})$. Then $*$ satisfies *9-*12 just in case $\leq_{\mathcal{B}}$ satisfies ≤ 5 - ≤ 8 .*

These approaches to iterated belief revision, of which Darwiche & Pearl (1997) is just one among many, have in common that the agent's beliefs are represented as a belief set plus an entrenchment ordering/system of spheres/plausibility ordering

both before and after the revision process. They also have in common that the new information is represented as a single sentence or a single proposition. The latter is also true for the approach by Jin & Thielscher (2007) discussed below, but different for what Rott (2009) calls “two-dimensional” belief revision operators (see Cantwell 1997 and Fermé & Rott 2004 and Rott 2007).

In “one-dimensional” belief revision the new information comes as a “naked” (Rott 2007) sentence or proposition, and it is the job of the various belief revision methods – lexicographic revision (Nayak 1994), natural revision (Boutilier 1996), irrevocable revision (Seegerberg 1998, Fermé 2000), irrefutable revision (Rott 2006), and others – to say where exactly in the new entrenchment ordering/system of spheres/plausibility ordering that sentence or proposition should be placed. In two-dimensional belief revision it is the new information itself that carries this information, at least to some extent. Here the new information does not merely say that the input sentence α should be accepted.² It additionally specifies how firmly α should be accepted by requiring α to be (at least) as entrenched as some “reference sentence” β so that the new information is of the form $\beta \leq \alpha$.

Let us return to our example. On Monday Sophia learns that the weather forecast for Tuesday and Wednesday predicts rain, δ . In one-dimensional belief revision she picks one of the iterated belief revision methods and revises her old belief set \mathcal{B} and entrenchment ordering $\leq_{\mathcal{B}}$ to obtain a new belief set $\mathcal{B} * \delta$ and entrenchment ordering $\leq_{\mathcal{B} * \delta}$. Different methods return different outputs, but on all of them she ends believing that it will rain on Tuesday, β . On Tuesday she sees that is sunny, $\neg\beta$. In one-dimensional belief revision Sophia proceeds as before. In two-dimensional belief revision Sophia does not merely receive the qualitative information $\neg\beta$, but the comparative information $\neg\gamma \leq \neg\beta$, indicating that she trusts her sight at least as much as the weather forecast. While there still are several belief revision methods to choose from (Rott 2009) – among others, reflecting the fact that Sophia can respect $\neg\gamma \leq \neg\beta$ by lowering the doxastic status of $\neg\gamma$ or by raising the doxastic status of $\neg\beta$ – the new information now is more specific and leaves less room to be filled by the revision method. It is then only a small step to equip Sophia with the quantitative, numerical information that $\neg\beta$ is entrenched to a specific degree. In that case the new information determines where $\neg\beta$ is located in the new entrenchment ordering. The revision method merely has to incorporate this into Sophia’s old beliefs in a coherent way. We will take this step in part II.

²For present purposes we can ignore the case of “non-prioritized” belief revision, where the new information need not be accepted. See Hansson & Fermé & Cantwell & Falappo (2001).

Before doing so let us briefly return to the qualitative approaches to iterated belief revision. As Jin & Thielscher (2007) point out the accounts by Darwiche & Pearl (1997) and others are still too permissive in the sense that beliefs are given up unnecessarily. As a consequence they additionally require the ideal epistemic agent to consider the new information β *independent* of a sentence α after revision by β if she did so before. While the idea behind Jin & Thielscher (2007)'s proposal seems to be correct, their actual proposal is too strong. The reason is that their notion of dependence is too strong in the sense that too many sentences are rendered independent. According to them a believed sentence α is independent of another sentence β if the believed sentence α is still believed after revision by the negation of the other sentence, $\neg\beta$. But I can receive new information $\neg\beta$, say, that the captain of my favorite football will not be fit for the match, whose negation β is positively relevant to, and so *not* independent of, a belief of mine α , say, that my favorite football team will win the match, without making me give up that belief of mine altogether. More generally, the ways two sentences can depend on each other are many and varied, and the qualitative and comparative notions of AGM belief revision theory seem to be too coarse-grained to capture those. Hild & Spohn (2008) argue axiomatically that, in order to adequately represent all dependencies, and to handle indefinitely iterated belief revisions, one has to go all the way to quantitative, numerical ranking functions.

Acknowledgments

I am grateful to Benjamin Bewersdorf, Jennifer Nagel, Hans Rott, Sherrilyn Roush, Wolfgang Spohn, Jonathan Weisberg, and an anonymous referee for helpful comments on earlier versions of these two papers. Robert Lehnert helped me with the design of the graphics. My research was in part supported by the German Research Foundation through its Emmy Noether program.

Short Biography

Franz Huber is an assistant professor in the Department of Philosophy at the University of Toronto. His research interests are in formal epistemology, general philosophy of science, and philosophical logic.

References

- Alchourrón, Carlos E. & Gärdenfors, Peter & Makinson, David (1985), On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic* **50**, 510-530.
- 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.
- Easwaran, Kenny (2011a): Bayesianism I: Introduction and Arguments in Favor. *Philosophy Compass* **6**, 312-320.
- Easwaran, Kenny (2011b): Bayesianism II: Applications and Criticisms. *Philosophy Compass* **6**, 321-332.
- Fermé, Eduardo (2000), Irrevocable Belief Revision and Epistemic Entrenchment, *Logic Journal of the IGPL* **8**, 645-652.
- Fermé, Eduardo & Rott, Hans (2004), Revision by Comparison, *Artificial Intelligence* **157**, 5-47.
- Gärdenfors, Peter (1988), *Knowledge in Flux. Modeling the Dynamics of Epistemic States*. Cambridge, MA: MIT Press.
- Gärdenfors, Peter & Makinson, David (1988), Revisions of Knowledge Systems Using Epistemic Entrenchment. *Proceedings of the 2nd conference on Theoretical aspects of reasoning about knowledge*. San Francisco: Morgan Kaufmann, 83-95.
- Gärdenfors, Peter & Rott, Hans (1995), Belief Revision. In D.M. Gabbay & C.J. Hogger & J.A. Robinson (eds.), *Epistemic and Temporal Reasoning. Handbook of Logic in Artificial Intelligence and Logic Programming: Volume 4*. Oxford: Clarendon Press, 35-132.
- Grove, Adam (1988), Two Modellings for Theory Change. *Journal of Philosophical Logic* **17**, 157-170.

- Hansson, Sven O. & Fermé, Eduardo & Cantwell, John & Falappa, Marcelo A. (2001), Credibility-Limited Revision, *Journal of Symbolic Logic* **66**, 1581-1596.
- Harper, William L. (1976), Rational Conceptual Change. *PSA 1976* **2**, 462-494.
- Hild, Matthias & Spohn, Wolfgang (2008), The Measurement of Ranks and the Laws of Iterated Contraction. *Artificial Intelligence* **172**, 1195-1218.
- Jin, Yi & Thielscher, Michael (2007), Iterated Belief Revision, Revised. *Artificial Intelligence* **171**, 1-18.
- Levi, Isaac (1977), Subjunctives, Dispositions and Chances. *Synthese* **34**, 423-455.
- Lewis, David K. (1973), *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Nayak, Abhaya C. (1994), Iterated Belief Change Based on Epistemic Entrenchment, *Erkenntnis* **41**, 353-390.
- Rott, Hans (2001), *Change, Choice, and Inference. A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford: Oxford University Press.
- 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 (2006), Revision by Comparison as a Unifying Framework: Severe Withdrawal, Irrevocable Revision and Irrefutable Revision, *Theoretical Computer Science* **355**, 228-242.
- Rott, Hans (2007), Two-dimensional Belief Change: An Advertisement. In G. Bonanno & J. Delgrande & J. Lang & H. Rott (eds.), *Formal Models of Belief Change in Rational Agents. Dagstuhl Seminar Proceedings* **07351**.
- Rott, Hans (2009), Shifting Priorities: Simple Representations for Twenty-seven Iterated Theory Change Operators. In D. Makinson & J. Malinowski & H. Wansing (eds.), *Towards Mathematical Philosophy. Trends in Logic* **28**. Dordrecht: Springer, 269-296.
- Seegerberg, Krister (1998), Irrevocable Belief Revision in Dynamic Doxastic Logic. *Notre Dame Journal of Formal Logic* **39**, 287-306.