Coherence and Conservatism in the Dynamics of Belief II: Iterated Belief Change Without Dispositional Coherence

Hans Rott

3 April 2002, 17:07

1 Introduction

Part I of this paper (Rott 1999) has argued that doxastic states should be represented as revision functions that take sequences of sentences (of arbitrary finite length) as arguments, and return belief sets (that is, sets of sentences closed under some background logic Cn) as a result. We have interpreted the well-known rationality postulates for belief change given by Alchourrón, Gärdenfors and Makinson (1985) and isolated three independent concepts of coherence that are (partially) encoded by them. We noticed that contrary to the common opinion, the classic theory expounded by these authors (and further developed by many more writers) does not offer a substantial account of the idea of *minimal changes* of belief.[1] This idea, which is also known as *conservatism*, embodies a *diachronic* conception of coherence. In the belief revision theory of Alchourrón, Gärdenfors and Makinson (henceforth, AGM), a much more important role is played by what I call *synchronic* coherence (or inferential coherence, or reflective equilibrium, roughly: logical consistency and closure of belief sets) and *dispositional* coherence. The latter concept finds expression in AGM's 'supplementary' revision postulates and is equivalent to the rationalizability of the revision behaviour by some very well-behaved choice function, viz., one that can be obtained by maximization with respect to a transitive preference ordering. Representation theorems showing that such equivalences do in fact obtain may justly be called the core of the AGM paradigm.[2]

[1] Also see Rott (2000).

[2] See the extensive overviews given in Gärdenfors (1988), Gärdenfors and Rott (1995), and Hansson (1999). In other work (Rott 1993, 1998, 2001), I have discussed how the supplementary AGM postulates can be systematically *weakened* by lessening the strictness of the conditions on "rational" choices. In this paper, I want to do without *any* kind of dispositional coherence.

In this paper I want to study the idea of conservatism, not conservatism with respect to beliefs (which actually plays a very modest role in existing belief change theories[3]) but conservatism with respect to belief change strategies. In order to eliminate an uncontrolled interference with the idea of dispositional coherence, I shall pay special attention to the case of *basic belief change* where neither the (logically strong) supplementary AGM postulates nor the (comparatively weak) AGM postulates concerning conservatism with respect to beliefs need to be satisfied.

The plan of the paper is as follows. We first recast the AGM postulates in the new setting of unary, iterated belief revision functions. Then we briefly recapitulate the theory of *basic entrenchment* (Rott 2001) that is suitable exactly for basic belief change. One-step belief change generated by basic entrenchment will then be combined with a natural and conservative method of revising entrenchment relations. We show that this method satisfies an additional postulate for iterated belief change, and that all basic iterated belief change functions satisfying this postulate can be represented as being generated by this conservative method of entrenchment revision. In the concluding discussion we compare our method with a number of qualitative approaches known from the relevant literature. It turns out that the basic conservative method generalizes a semantic approach advocated by Boutilier (1993, 1996), and that our method – though not relying on the power of dispositional coherence – unfortunately has the same fundamental problem as Boutilier's approach. Our diagnosis in conclusion is that unmediated conservativity is no good. Conservative revision suffers from a problem of *temporal* incoherence – which points to a fourth kind of coherence that is in need of further investigation.

2 Elements of iterated belief change

In part I of this paper (Rott 1999) we went through a longish reflection upon the proper representation of belief states. We were led to the conclusion that belief states should be represented by one-place iterated belief revision functions. In the previous paper, various concepts of coherence were discussed with reference to one-step functions. It is now time to develop the precise tools for a theory that takes iterated belief revision functions as primitive.

Some preliminaries. We work in a propositional language with the usual connectives. $\mathcal{L}$ is the set of sentences of this language (sometimes identified with the language). The propositional constants $\top$ ("the truth") and $\bot$ ("the falsity") are in $\mathcal{L}$. The variables ϕ, ψ, χ, etc. range over sentences from $\mathcal{L}$. We will often talk about sequences of $\mathcal{L}$-sentences. Instead of the cumbersome $\langle \phi_1, \ldots, \phi_n \rangle$ and

[3] Compare Rott (2000).

$*(\langle\phi_1,\ldots,\phi_n\rangle)$, we shall always use capital Greek letters and write Φ and $K*\Phi$, with the understanding that $\Phi = \langle\phi_1,\ldots,\phi_n\rangle$. For the degenerate case of the empty input sequence $\mathbf{0}$, when no revision takes place at all, $K*\Phi = K*\mathbf{0}$ is also called 'K'. A sequence $\langle\phi\rangle$ of length 1 is identified with the sentence ϕ. The capital Greek letter Ψ will stand for $\langle\psi_1,\ldots,\psi_m\rangle$. Concatenation of sequences is denoted by a dot, $\cdot$, so $\Phi\cdot\Psi$ for instance denotes $\langle\phi_1,\ldots,\phi_n,\psi_1,\ldots,\psi_m\rangle$. We presume that the language $\mathcal{L}$ is goverened by a compact Tarskian (i.e., refexive, monotonic and idempotent) consequence operation Cn. Where convenient, we write $\phi\vdash\psi$ for $\psi\in Cn(\phi)$, and we use the abbreviation $K+\phi$ for the logical *expansion* $Cn(K\cup\{\phi\})$ of K by ϕ. Notice that the '+' in '$K+\phi$' does not stand for a dynamic operation of belief change, but for a purely logical notion.

A *unary iterated belief revision function* $*$ assigns, for each finite sequence $\langle\phi_1,\ldots,\phi_n\rangle$ of input sentences, the belief set (i.e., the logically closed set of sentences) obtained after successively revising the current belief state by ϕ_1, by ϕ_2, ..., and finally by ϕ_n. Formally, the revision function $*$ is a function with domain $\mathcal{L}^\omega$ and range $\mathbb{K}$. Functions of this signature will be taken to represent *belief states*. We associate with a belief state $*$ a belief set $K = *(\mathbf{0})$ that is obtained by a vacuous "revision" with the empty sequence $\mathbf{0} = \langle\ \rangle$.

Although the AGM postulates have been devised for unary, one-step belief revision, it is straightforward to translate them into the context of iterated belief revision. We just have to be aware of the fact that an iterated belief revision function gives rise to infinitely (but countably) many one-step revision functions $*_\Phi$ for each belief set $K*\Phi$, definable from $*$ by the equation $(K*\Phi)*_\Phi\psi = K*(\Phi\cdot\psi)$.

Appropriate collections of the AGM-style postulates are now required to hold not only for the belief set K at the beginning of a 'doxastic walk', but also for any $K*\Phi$ reachable from K through a sequence of revisions by $\Phi = \langle\phi_1,\ldots,\phi_n\rangle$. As we explained in Part I of this paper, however, it is wrong to think of the AGM postulates as implicitly quantifying over all belief sets. Belief sets come into play only indirectly as the results of feeding the belief state $*$ with some sequence of input sentences. The right reading of the postulates is that a belief state, represented by an iterated revision function, may be called rational (in some specific sense) if and only if it – or more exactly, the revision function representing it – satisfies (some specific selection of) the postulates that we are now going to formulate. The notation '$K*\Phi$' is just an alternative way of referring to $*(\Phi)$, a way that facilitates comparison with the well-known AGM theory. It should not be mistaken to indicate that the postulates talk about two-place revision functions, i.e., functions taking as arguments pairs consisting of a belief set and a sentence.

From now on, a unary iterated function $*$ is supposed to satisfy at least the *basic set* of postulates for belief revision. These postulates derive from AGM (we keep AGM's numbering), but they are now presented in a format that is appropriate for unary iterated belief revision functions:

$(*1^\omega)$	$K * \Phi = Cn(K * \Phi)$	*(Closure)*
$(*2^\omega)$	$\psi \in K * (\Phi \cdot \psi)$	*(Success)*
$(*5^\omega)$	If $\psi \nvdash \bot$, then $K * (\Phi \cdot \psi) \nvdash \bot$	*(Consistency)*
$(*6^\omega)$	If $\psi \dashv\vdash \chi$, then $K * (\Phi \cdot \psi) = K * (\Phi \cdot \chi)$	*(Extensionality)*

The AGM-style postulates of *minimal change in the consistent case* are distinguished from the basic postulates as being less mandatory.[4]

$(*3^\omega)$	$K * (\Phi \cdot \psi) \subseteq (K * \Phi) + \psi$	*(Expansion)*
$(*4^\omega)$	If $\neg\psi \notin K * \Phi$, then $K * \Phi \subseteq K * (\Phi \cdot \psi)$	*(Preservation)*

In the present context $(*3^\omega)$ and $(*4^\omega)$ do not relativize the iterated revision function $*$ to a pregiven belief set K. As explained above, the revision function $*$ is primitive, and the associated belief set is derived from $*$ by means of the definition $K = *(\mathbf{0})$.

The postulates for dispositional coherence – strong postulates characteristic for AGM – complete our list.

$(*7^\omega)$ $\quad K * (\Phi \cdot (\psi \wedge \chi)) \subseteq (K * (\Phi \cdot \psi)) + \chi$

$(*8^\omega)$ $\quad$ If $\neg\chi \notin K * (\Phi \cdot \psi)$, then $K * (\Phi \cdot \psi) \subseteq K * (\Phi \cdot (\psi \wedge \chi))$

For explanation and motivation, we refer the reader to part I of this paper. Let us call an iterated revision function $*$ that satisfies postulates $(*1^\omega)$, $(*2^\omega)$, $(*5^\omega)$ and $(*6^\omega)$ a *basic iterative revision function*. If $*$ is basic and in addition satisfies $(*3^\omega)$ and $(*4^\omega)$, then we call it *c-conservative* ('c' for 'consistent case'); if it is basic and in addition satisfies $(*7^\omega)$ and $(*8^\omega)$, then we call it *dispositional*. If $*$ satisfies postulates $(*1^\omega)$ – $(*8^\omega)$ we call it an *iterative AGM revision function*.[5] Clearly then we can get a one-step basic (c-conservative, dispositional or AGM) revision function from an iterative basic (respectively, c-conservative, dispositional or AGM) revision function by restricting the latter's domain to $\mathcal{L}$. We then validate the original AGM postulates $(*1)$ – $(*8)$.

It is important to note that we do *not* have a constraint like

[4] AGM themselves are more demanding and call $(*3)$ and $(*4)$ "basic" as well.

[5] Lehmann (1995) who works in essentially the same framework as we do, also presents a postulate that is supposed to generalize the AGM axioms $(*7)$ and $(*8)$ to the context of iterated belief changes. One form of it is

(I6′) $\quad$ If $\neg\chi \notin K * (\Phi \cdot \psi)$, then $K * (\Phi \cdot \psi \cdot \chi \cdot \Psi) = K * (\Phi \cdot \psi \wedge \chi \cdot \Psi)$

This clearly implies $(*7^\omega)$ and $(*8^\omega)$, but it is *much* stronger than these postulates, since it says that in the consistent case the effects of updating with ψ and χ successively and the effects of updating with $\psi \wedge \chi$ are the same not only as regards the beliefs, but also as regards the dispositions to further change the beliefs.

(Ahistoricity) If $K * \Phi = K * \Psi$ then $K * (\Phi \cdot \chi) = K * (\Psi \cdot \chi)$

It is a basic intuitive presumption of the present paper that identical belief sets having different "histories" will in general be revised differently at later revision stages. In Section 5 of part I of this paper, I have argued that "in reality" we don't revise belief sets, but we revise the initial *belief state* $*$ into the *belief state* $*_\Phi$ which is defined by $*_\Phi(\Psi) = *(\Phi \cdot \Psi)$.[6] Below (at the end of Section 6) we shall discuss examples that make clear how this point gets reflected in our model.

The sequential generalization of the AGM postulates brings out the fundamental difference between the pairs $(*3)/(*4)$ and $(*7)/(*8)$ more clearly than the original one-step version. While the former pair relates prior and posterior belief states (revisions by sequences of different lengths), the latter pair relates two posterior belief states which are occasioned by different but logically related inputs (revisions by sequences of the same length). As in the one-step case, $(*3^\omega)$ and $(*4^\omega)$ say nothing at all about the interesting case where the input ψ is inconsistent with the belief set $K * \Phi$.[7] And as in the one-step case, $(*3^\omega)$ and $(*4^\omega)$ are the only postulates saying anything at all about minimal changes of belief sets – even if only for the case of consistent revisions.

We thus find ourselves in an unexpected predicament. So far we do not have a clue what to make of the concept of diachronic coherence or conservatism if revisions are belief-contravening, not even on the simplifying assumption that belief states can sensibly be represented by belief sets. It will turn out, rather surprisingly, that in order to get a hold on these concepts we need to have a constraint that does not only transfer an AGM postulate into a format suitable for iterated belief revision functions, but deals itself with genuinely iterated changes.

In order to find our way out of the predicament, we take a detour via entrenchment relations. We use the theory of basic entrenchment in the context of basic one-step revision functions, as developed in Rott (2001).

3 Basic entrenchment

In this section we give a brief summary of the concepts and the most relevant results concerning *basic entrenchment relations* (Rott 2001) which show that the well-known theory of epistemic entrenchment due to Gärdenfors and Makinson

[6] On the simplifying assumption that identifies belief states with belief sets, the concept of *synchronic coherence* is as easy to formalize as the concept of *dispositional coherence*. The former is captured by $(*1^\omega)$ and $(*5^\omega)$, the latter by $(*7^\omega)$ and $(*8^\omega)$. We refrain from discussing how to capture these coherence constraints in the formalism that identifies belief states with revision functions.

[7] Nor do $(*7^\omega)$ and $(*8^\omega)$, to be sure, but it is simply not the point of this pair of postulates to relate $K * \Phi$ to $(K * \Phi) * \psi$, or any other pair of prior and posterior states.

(1988) is applicable in a much wider context than originally conceived. Out of the eight AGM axioms, only four are necessary to make the entrenchment construction work. We first develop this theory in the context of one-step revision functions and generalize it for the case of iterated functions only later.

An entrenchment relation is a binary relation $\leq$ over the sentences in $\mathcal{L}$. More specifically, any relation over $\mathcal{L}$ that satisfies the following four conditions will be called a *basic entrenchment relation*.

(Reflexivity)	$\phi \leq \phi$
(Extensionality)	If $\phi \dashv\vdash \psi$ then :
	$\phi \leq \chi$ iff $\psi \leq \chi$, and $\chi \leq \phi$ iff $\chi \leq \psi$
(Choice)	$\phi \wedge \psi \leq \chi$ iff $\phi \leq \psi \wedge \chi$ or $\psi \leq \phi \wedge \chi$
(Maximality)	If $\top \leq \phi$ then $\vdash \phi$

Basic relations over $\mathcal{L}$ that in addition satisfy the following two conditions will be called *faithful* with respect to the belief set K:

(K-Minimality)	If $\phi \notin K$ then $\phi \leq \psi$
(K-Representation)	If $\phi \in K$ and $\phi \leq \psi$ then $\psi \in K$

Faithful relations that in addition satisfy Transitivity are called *standard entrenchment relations.*

Every entrenchment relation $\leq$ vacuously satisfies both K-Minimality and K-Representation with respect to the inconsistent belief set $K = \mathcal{L}$. If there is a consistent belief set K with respect to which $\leq$ satisfies K-Minimality and K-Representation, then K is uniquely determined as $K = \{\phi : \bot < \phi\}$ (see Rott 2001). By convention, we define for every entrenchment relation $\leq$ the associated belief set as follows:

(Def $K_{\leq}$) $\qquad K_{\leq} = \{\phi : \bot < \phi\}$

The postulate Choice receives its name from its interpretation as a central feature of entrenchments in a framework using "syntactic choice functions" (Rott 2001, Chapters 7 and 8). The left-hand side and the right-hand side of Choice are both ways of expressing that either ϕ or ψ is given up in a situation in which at least one of ϕ, ψ and χ needs to be given up. While Reflexivity, Extensionality and Choice are structural properties, Maximality, K-Minimality and K-Representation concern the limiting cases of tautologies and non-beliefs.

Entrenchment relations are used for constructing changes of belief. For our purposes, however, it is best to first look at the converse, *re*constructive interpretation of entrenchment. Given some unary, one-step revision function $*$, an entrenchment relation $\leq$ can be *retrieved from* $*$ by means of the following definition:[8]

$$(\text{Def} \leq \text{from} *) \qquad \phi \leq \psi \quad \text{iff} \quad \phi \notin K * \neg(\phi \wedge \psi) \ \text{ or } \ \vdash \psi$$

This condition expresses what we might call the *meaning of entrenchment*. It says, in the principal case, that ϕ is not more entrenched than ψ in an agent's belief state if and only if the revision of the belief state occasioned by the information that not both ϕ and ψ are true leads to a state in which ϕ is eliminated. In order to appreciate the import of this concept, it is necessary to understand that *all basic revision functions* can be represented as revision functions based on an underlying relation of entrenchment. Proofs for the observations of this section are given in Rott (2001).

Observation 1. If $* : \mathcal{L} \rightarrow \mathbb{K}$ is a basic one-step revision function, i.e., $*$ satisfies $(*1)$, $(*2)$, $(*5)$ and $(*6)$, and $\leq$ is the entrenchment relation over $\mathcal{L}$ retrieved from $*$ by means of (Def $\leq$ from $*$), then $*$ can be represented with the help of $\leq$ in the following way:

$$(\text{Def} * \text{from} \leq) \qquad \psi \in K * \phi \quad \text{iff} \quad \neg\phi < \phi \supset \psi \ \text{ or } \ \vdash \neg\phi$$

Here $<$ is the asymmetric part of $\leq$. The principal part of condition (Def $*$ from $\leq$) says that ψ is in $K * \phi$ if the material conditional $\phi \supset \psi$ is strictly more entrenched than $\neg\phi$, i.e., the negation of the input sentence. If a (one-place) revision function $*$ is defined from a relation $\leq$ in $\mathcal{L}$ with the help of (Def $*$ from $\leq$), we say that $*$ is *based* on $\leq$ (or that $\leq$ *determines* $*$). If $\leq$ is retrieved from a basic revision function $*$ quite a number of properties of $\leq$ can be derived:

Observation 2. (a) If $*$ is a basic (one-step) revision function satisfying $(*1)$, $(*2)$, $(*5)$ and $(*6)$, then the relation $\leq$ retrieved from $*$ is a basic entrenchment relation.

(b) If $*$ in addition satisfies $(*3)$, then $\leq$ satisfies (K-Minimality).

(c) If $*$ in addition satisfies $(*4)$, then $\leq$ satisfies (K-Representation).

(d) If $*$ in addition satisfies $(*7)$ and $(*8)$, then $\leq$ is transitive.

In order to secure transitivity for the entrenchment relation retrieved from a revision function $*$, the full power of postulates $(*7)$ and $(*8)$ is badly needed.

Although the conditions for basic entrenchment are definitely non-trivial, they do not guarantee acyclicity. But basic entrenchment relations satisfy a number of other important properties.

[8] Essentially the same definition as applied to belief *contraction* is due to Gärdenfors and Makinson (1988); it was transferred to belief *revision* by Rott (1991b), amongst others.

Lemma 3. Let $\leq$ satisfy Reflexivity, Extensionality and Choice. Then it also satisfies the following properties:

(Conjunctiveness)	$\phi \leq \psi$ iff $\phi \leq \phi \wedge \psi$
(Conditionalization)	$\phi \leq \psi$ iff $\phi \leq \phi \supset \psi$
(Connectedness)	$\phi \leq \psi$ or $\psi \leq \phi$
(GM-Dominance)	If $\phi \vdash \psi$ then $\phi \leq \psi$
(GM-Conjunctiveness)	$\phi \leq \phi \wedge \psi$ or $\psi \leq \phi \wedge \psi$

If $\leq$ in addition satisfies Maximality, then it also satisfies

(GM-Maximality)	If $\psi \leq \phi$ for all ψ then $\vdash \phi$

If $\leq$ in addition satisfies K-Minimality and K-Representation, then it also satisfies

(GM-Minimality)	If $K \neq K_\perp$ and $\phi \leq \psi$ for all ψ then $\phi \notin K$

The relations considered in Lemma 3 are required to meet far less demanding conditions than the entrenchment relations of Gärdenfors and Makinson (1988) which can only be retrieved from revision functions satisfying the full set of AGM postulates for revision. While basic entrenchment is in general not transitive, standard entrenchment is. In fact Gärdenfors and Makinson characterize their *GM-entrenchment relations* by the set consisting of GM-Dominance, GM-Conjunctiveness, GM-Maximality, GM-Minimality plus Transitivity. Thus Transitivity is the only GM-feature that basic entrenchments miss. Conversely, we have:

Observation 4. GM-entrenchment relations satisfy all conditions of basic entrenchments.

The next result says that if a basic revision function $*$ is determined by a basic entrenchment relation, then this entrenchment relation must be exactly the one which is retrievable from $*$.

Observation 5. Let the entrenchment relation $\leq$ satisfy Reflexivity, Extensionality and Choice, and let $*$ be based on $\leq$. Then $\leq$ can be retrieved from $*$ with the help of (Def $\leq$ from $*$).

The final result reveals that the constraints for basic entrenchment relations make sure that the revisions based on them satisfy the basic postulates for revison functions.

Observation 6. (a) If $\leq$ is a basic entrenchment relation satisfying Reflexivity, Extensionality, Choice and Maximality, then the revision function $*$ based on $\leq$ satisfies $(*1)$, $(*2)$, $(*5)$ and $(*6)$.

(b) If $\leq$ in addition satisfies K-Minimality, then $*$ also satisfies $(*3)$.

(c) If $\leq$ in addition satisfies K-Representation, then $*$ also satisfies $(*4)$.

(d) If $\leq$ in addition satisfies Transitivity, then $*$ also satisfies $(*7)$ and $(*8)$.

4 The use of basic entrenchment in iterated belief change

The Minimality and Representation conditions tie an entrenchment relation to a unique belief set K. One might argue that they are in fact essential the meaning of "entrenchment". A sentence that is not even believed to be true is not entrenched in the belief state at all. All non-beliefs, and no beliefs, are assigned the least degree of entrenchment. The entrenchment relation must be faithful to K.

The conditions of Minimality and Representation create a problem which has led some researchers to think that the AGM approach to belief revision is completely impotent as regards iterated belief changes. Consider the following argument to the effect that relations of epistemic entrenchment *cannot* be applied in iterated changes of belief.

> All and only sentences that are not believed to be true by the agent receive a minimal degree of entrenchment, i.e., entrenchment relations have to satisfy the conditions of K-Minimality and K-Representation. Thus entrenchment relations are intrinsically dependent on the current set of the agent's beliefs.
> So if the agent's belief set changes, the old entrenchment relation that was tailored to fit the old belief set does not fit the new belief set any more, and *therefore* it becomes useless.
> Hence we are left without any guidance for further changes of the revised belief set.

We will show that this argument is faulty. It breaks down at the point marked by the italicized "therefore" in the text. The old entrenchment relation can still be of use, even after the belief set has changed. This conforms to an idea sketched in Section 3 of part I of this paper: Entrenchment relations can be regarded as components of representations of mental states – components that may well be assumed to persist longer than (large portions of) the believed propositions themselves. The values attached to beliefs, or the preferences between beliefs, may remain the same even though the beliefs themselves are continuously subject to revision. The problem we are facing is that we need to reconcile this idea with the seemingly incompatible idea expressed by the Minimality and Representation conditions.

In order to solve this problem for belief revision , we will apply the *prior* entrenchment relation to the *posterior* belief set. But we also *restrict* the application of the prior entrenchment relation and make sure that the posterior entrenchment of posterior non-beliefs is taken care of by a separate clause in such a way that the conditions of Minimality and Representation are satisfied. In a loose sense, we essentially keep our belief revision strategy (encoded in the entrenchment relation) while we change our beliefs. But of course, this is not strictly speaking true, since taking care of Minimality and Representation does involve some non-negligible changes of the entrenchment relation.

The strategy of keeping the *entrenchments* of beliefs "as much as possible" (i.e., as much as allowed by Minimality and Representation) even across varying belief sets reflects *a conservative attitude*. This attitude is different – and in fact independent – from the kind of conservatism often advocated in epistemology and belief revision that advises us to retain as many *beliefs* as possible. In Rott (2000), I have argued that surprisingly, the latter attitude has not played any major part in the belief revision literature so far. In the next section, we will explore the consequences of adopting the conservative strategy with respect to one's entrenchment relation.

Before doing that, let us take stock of our findings so far. Entrenchment relations stand in a one-to-one correspondence to (unary) one-step revision functions. We saw that only very few postulates are required in order to reconstruct such a revision function as a function based on an entrenchment relation, viz., ($*1$), ($*2$), ($*5$) and ($*6$). Basic entrenchment relations can be retrieved from such revision functions, and the method of retrieval can easily be extended to the context of iterated revisions. Given a belief state $*$, every sequence of inputs Φ defines an entrenchment relation:

(Def$^\omega$ $\leq$ from $*$) $\quad\quad \psi \leq_\Phi \chi \quad \text{iff} \quad \psi \notin K * (\Phi \cdot \neg(\psi \wedge \chi)) \ \text{ or } \ \vdash \chi$

An iterated belief revision function thus yields an entrenchment relation $\leq_\Phi$ for each belief set $K * \Phi$ that can be reached by revising the initial belief set K through a sequence Φ of potential inputs. Note that $K * \Phi = K * \Psi$ does not imply $\leq_\Phi = \leq_\Psi$.

The theorems established in Section 3 for the one-step case carry over to the iterated case. Under quite weak conditions, the one-step revision function that is to be applied after an input sequence Φ has been accommodated can be represented as a revision function based on $\leq_\Phi$:

(Def$^\omega$ $*$ from $\leq$) $\qquad \chi \in K * (\Phi \cdot \psi)$ iff $\neg\psi <_\Phi \psi \supset \chi$ or $\vdash \neg\psi$

Instead of $K * (\Phi \cdot \psi)$, we can also write $(K * \Phi) *_\Phi \psi$. Tor each sequence Φ, the correspondences between the one-step revision function $*_\Phi$ and the entrenchment relation $\leq_\Phi$ are precisely the same as those described in Section 3.

Another way of looking at the development of epistemic states is available by defining a revision function $* : \mathcal{L}^\omega \rightarrow \mathbb{E}$. While a single entrenchment relation $\leq$ in $\mathbb{E}$ encodes much more information than a single belief set K in $\mathbb{K}$, the function $*$ is not more informative than the function $*$, thanks to the retrievability of entrenchments from (one-step) revisions.

It remains to clarify the way how a given entrenchment relation can be utilized for *iterated* changes of belief. Two different ways of achieving this are possible. The first idea is to use the current entrenchment relation for the construction of the new belief set $K * \phi$ with the help of (Def $*$ from $\leq$), and then use $K * \phi$ for the transformation of $\leq$ into the new entrenchment relation $\leq' = \leq^*_\phi$. It will turn out that this idea is indeed crucial for conservative belief change, but it cannot be generalized to other approaches. The second idea is to have a principled method of directly transforming the prior entrenchment relation $\leq$ into a new entrenchment relation $\leq' = \leq^*_\phi$, and only then obtaining the new belief set by putting $K * \phi = K_{\leq^*_\phi}$. Since an entrenchment relation encodes more information than a belief set, this method is generally applicable, even if it is, in some case, less easily accessible than the former method.

But what is a "method of directly transforming entrenchment relations"? It is a constructive suggestion for how to change entrenchment relations in the light of new evidence, more formally, where $\mathbb{E}$ is the set of all entrenchment relations, it is a function

$$\circledast : (\mathbb{E} \times \mathcal{L}) \rightarrow \mathbb{E}$$

that takes a prior entrenchment relation $\leq$ and an input sentence ϕ and returns a posterior entrenchment relation $\leq^*_\phi$ with $\perp <^*_\phi \phi$ (the latter meaning that ϕ is accepted in the posterior belief set $K_{\leq^*_\phi}$).

Notice that while we have always insisted on the historical sensitivity of the revision of belief sets, we conceive of the revision of entrenchment relations as functionally determined by the input sentence. Given a specific method of entrenchment revision, everything that might be relevant to a specific act of revision is encoded in the prior entrenchment relation itself.

Let $\mathbb{K}$ be the set of all belief sets. From now on we want to focus on travels in the 'doxastic space' $\mathbb{K}$ that are *determined by* (or *generated by* or *based on*) a basic entrenchment relation $\leq$ (the initial condition as it were) and a certain method $\circledast$ for revising entrenchment relations.

If each one-step revision function is based on the then current epistemic entrenchment relation, a revision operation on entrenchments determines an *iterated* revision function for the original belief set $K = K_{\leq}$. All we need to start from is a single entrenchment relation. Therefore, if a specific method for entrenchment revision is agreed upon, we can in fact identify a *belief state* with a single entrenchment relation (rather then with an iterated revision function $*$).

Instead of the cumbersome $(\ldots((\leq^*_{\phi_1})^*_{\phi_2})\ldots)^*_{\phi_n}$, we will again use Greek capital letters and write $\leq^*_\Phi$. This is the entrenchment relation that results from repeated entrenchment revisions by a sequence $\phi_1, \phi_2, \ldots, \phi_n$ of input sentences, always using the same method of entrenchment revision. For the degenerate case $\Phi = \mathbf{0}$, $\leq^*_\Phi$ is always identified with $\leq$.

It is now easy to respresent the way how a single initial entrenchment relation and a given method of revisions entrenchment relations determines an iterated revision function. An iterated revision function $* : L^\omega \to \mathbb{K}$ is *generated by* (or *determined by* or *based on*) an entrenchment relation $\leq$ and a method $\circledast$ of entrenchment revision if and only if for all $\Phi = \langle\phi_1, \ldots, \phi_n\rangle \neq \mathbf{0}$ with consistent ϕ_n,

$$*(\Phi) = K_{\leq^*_\Phi}$$

where $\leq^*_\mathbf{0} = \leq$ and $\leq^*_{\Phi\cdot\psi}$ is defined from $\leq^*_\Phi$ by this method $\circledast$ of entrenchment revision, that is, $\leq^*_{\Phi\cdot\psi} = \circledast(\leq^*_\Phi, \psi)$.

For the case $\Phi = \mathbf{0}$ we require that $\leq\, = \,\leq^*_\mathbf{0}$ satisfies K-Minimality and K-Representation with respect to $K = *(\mathbf{0})$, i.e., that either $*(\mathbf{0}) = K_{\leq}$ (if $\leq$ satisfies $\perp$-Minimality and $\perp$-Transitivity) or $*(\mathbf{0}) = \mathcal{L}$. For the case $\Phi = \langle\phi_1, \ldots, \phi_n\rangle$ with inconsistent ϕ_n, the identity $*(\Phi) = \mathcal{L}$ has to hold by $(*5^\omega)$, so $\leq^*_\Phi$ satisfies K-Minimality and K-representation with respect to $K * \Phi = *(\Phi)$ anyway, and we do not require that $*(\Phi) = K_{\leq^*_\Phi}$.

But there is also a second idea that we want to heed: Our theory of basic entrenchment. Therefore, we require the identity (Def $*$ from $\leq$) to hold as well.

So we have two ways of finding the new belief set: (Def $*$ from $\leq$) constructs $*(\Phi\cdot\psi)$ from the prior entrenchment relation $\leq^*_\Phi$, while (Def $K_{\leq}$) constructs $*(\Phi\cdot\psi)$ from the posterior entrenchment relation $(\leq^*_\Phi)^*_\psi$. We want to respect both ideas and let their results coincide. The situation is illustrated graphically in Fig. 1. It is therefore necessary that $\perp <^*_\phi \psi$ if and only if $\neg\phi < \phi \supset \psi$, for all ψ. In order to avoid an unilluminating tinkering with limiting cases, we require this identity to hold only for consistent ϕ. So equivalently, if $\neg\phi \in Cn(\emptyset)$, then we should have:

$$\psi \leq^*_\phi \perp \text{ if and only if } \phi \supset \psi \leq \neg\phi$$

Let us call this equivalence the *Triangle property* because it lays down that both ways of constructing $K * \phi$ must coincide. It specifies an important restriction on methods for entrenchment revision. All methods of entrenchment revision we

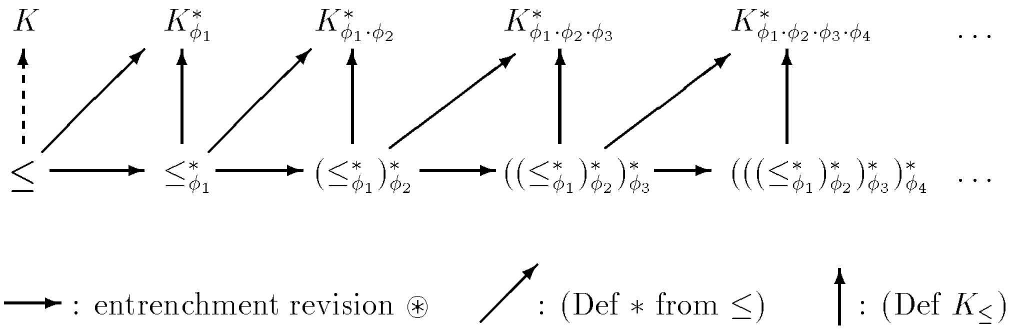

Figure 1: *Generating an iterated belief revision function* $* : \mathcal{L}^\omega \to \mathbb{K}$ *from an initial entrenchment relation* $\leq$ *and an entrenchment revision function* $\circledast : \mathbb{E} \times \mathcal{L} \to \mathbb{E}$

are going to consider satisfy the Triangle property.[9]

5 Conservatism

Intuitively, entrenchment relations do not and should not stay the same through a series of revisions.[10] They should be affected by experiences, they should be relative to the current belief set, and this relativity is formally nicely captured by K-Minimality and K-Representation. Thus, as beliefs get revised, entrenchment relations *need* to be revised as well. But there is reason to require that revisions of such "local", "theory-relative" entrenchment operations be conservative. So far we have not even the beginnings of a general method of capturing the conservativity of *changes of belief sets* that also works for inputs inconsistent with the current beliefs. (For the consistent case we have $(*3^\omega)$ and $(*4^\omega)$.) Fortunately, we can find such a method when we turn to *changes of entrenchment relations*, and it turns out that this method is very simple indeed.

As already mentioned, the basic idea is to use essentially the old entrenchment relation for the revised belief set, while at the same time seeing to it that the conditions of Minimality and Representation are satisfied with respect to the new belief set. We need not consider iterations in order to make this idea more precise. Let X be an arbitrary belief set (i.e., a logically closed set of sentences), let $\leq$ be

[9]The triangle property is similar in spirit to Boutilier's (1996, pp. 271–272) "Basic Requirement" on revision functions as applied to models.

[10]For models representing a different view of the matter, see the Section 8.1.

an entrenchment relation and let $\leq_X$ denote a new entrenchment relation that results from *relativizing* $\leq$ *to* X in the following way:

$$\psi \leq_X \chi \;\text{ iff }\; \begin{cases} \psi \notin X & \text{or} \\ \chi \in X \text{ and } \psi \leq \chi & \end{cases}$$

The relativized relation $\leq_X$ is identical with $\leq$ if either $\leq$ satisfies (X-Minimality) and (X-Representation) to begin with, or if X is inconsistent. Relativization to X makes the entrenchment relation faithful to X, but otherwise it preserves all the relevant structural properties:

Observation 7. Let X be a belief set and $\leq$ a relation over $\mathcal{L}$. If $\leq$ satisfies Reflexivity (or respectively, Extensionality, Choice, Maximality, Transitivity), so does the relativized relation $\leq_X$. Moreover, $\leq_X$ satisfies X-Minimality and X-Representation.

Given a basic (or given a standard) entrenchment relation $\leq$, the relativized relation $\leq_X$ is a basic (or, respectively, standard) entrenchment relation that can be used to determine a revision function which will satisfy $(*1)$ – $(*6)$ (or, respectively, $(*1)$ – $(*8)$), where $(*3)$ and $(*4)$ mention X rather than K. The scheme above can be applied to any arbitrary belief set X and yields a relation of entrenchment for X under preservation of the relevant structural features. In particular, then, this method can be applied to the revised belief sets themselves![11] Having defined a simple method of relativizing an entrenchment relation in such a way that it is faithful with respect to a given set, we now put, in a second step, $X = K * \phi$. We can now escape from the defeatist argument of the previous section and define the process of *conservative entrenchment revision* by putting $\leq_\phi^{*C} = \leq_{K*\phi}$. After revising the belief state with ϕ, the old entrenchment function gets relativized to the new belief set $K * \phi$.

$$\psi \leq_\phi^{*C} \chi \;\text{ iff }\; \begin{cases} \psi \notin K * \phi & \textit{or} \\ \chi \in K * \phi \text{ and } \psi \leq \chi & \end{cases}$$

Surely this is as simple and conservative as an entrenchment revision can possibly be. We know that $\leq_\phi^{*C}$ no longer satisfies Minimality and Representation with respect to K, but with respect to the revised belief set $K * \phi$ instead. *Within* the new belief set $K * \phi$, however, we keep on using the old entrenchment relation $\leq$.

Now we assume in addition that the (unary and one-step) revision function $*$ taking K to $K * \phi$ is itself based on the entrenchment relation $\leq$. The most intuitive way of performing conservative entrenchment revision takes two steps.

[11] Or to contracted belief sets, for that matter. See footnote 13.

First, $\leq$ is used as a basis to construct the new belief set $K * \phi$, and second, this new belief set $K * \phi$ is used to relativize $\leq$ and turn it into $\leq_\phi^{*C}$.

Notational conventions. Instead of the cumbersome $(\ldots((\leq_{\phi_1}^{*C})_{\phi_2}^{*C})\ldots)_{\phi_n}^{*C}$, we will write $\leq_\Phi^{*C}$, with the understanding that $\Phi = \langle \phi_1, \ldots, \phi_n \rangle$. This is the entrenchment relation that results from repeated conservative entrenchment revisions by a sequence $\phi_1, \phi_2, \ldots, \phi_n$ of input sentences. It must be distinguished carefully from the more abstractly defined relations $\leq_\Phi$ and $\leq_\Phi^*$ of the previous section that do not refer to the specific approach using conservative revisions.

Let us now combine (Def$^\omega *$ from $\leq$) with the recipe of conservative entrenchment revision. We are finally in a position to give the formal definition of *conservative iterated belief revision* based on entrenchment relations. Let $\leq$ be the basic entrenchment relation from which we start. We simultaneously define the iterated revision function $* : \mathcal{L}^\omega \to \mathbb{K}$ and the entrenchment relations $\leq_\Phi^{*C}$ associated with each revised belief set $*(\Phi)$. The definition is by induction. For the sequence $\mathbf{0}$ with length zero we define the belief set

$$*(\mathbf{0}) = K = \begin{cases} K_\leq & \text{, if } \leq \text{ is faithful to } K_\leq \\ \mathcal{L} & \text{, otherwise}^{12} \end{cases}$$

and the entrenchment relation

$$\leq_{\mathbf{0}}^{*C} = \leq$$

For the induction step, we assume that we have the belief set $*(\Phi)$ and the entrenchment relation $\leq_\Phi^{*C}$ that result after accommodating the input sequence Φ of length n. For the sequence $\Phi \cdot \psi$ of length $n+1$, we define the belief set

$$*(\Phi \cdot \psi) = \{\chi : \neg\psi <_\Phi^{*C} \psi \supset \chi \text{ or } \vdash \neg\psi\}$$

and the entrenchment relation

$$\chi \leq_{\Phi\cdot\psi}^{*C} \xi \quad \text{iff} \quad \begin{cases} \chi \notin *(\Phi \cdot \psi) & \textit{or} \\ \xi \in *(\Phi \cdot \psi) \text{ and } \chi \leq_\Phi^{*C} \xi \end{cases}$$

If this procedure is followed, we say that $\leq$ *conservatively determines* the iterated revision function $*$. It is evident that this method combines the Gärdenfors-Makinson recipe ($*$ from $\leq$) for constructing entrenchment-based revisions with the conservative revision revision strategy (relativization to the new belief set) of entrenchment relations.

Conservative entrenchment revision can be characterized without any reference to belief sets at all. This is methodologically sound strategy, but unfortunately, the transformed condition loses a lot of the simplicity and intuitive appeal of the original idea. For any input sentence ϕ that is not contradictory we get

[12] This guarantees that $\leq$ is faithful to $*(\mathbf{0})$. With the lower line, we are able to describe a doxastic travel as an escape from "epistemic hell". Compare page 30 below.

$$\text{(CER)} \qquad \psi \leq^{*C}_{\phi} \chi \;\text{ iff }\; \begin{cases} \phi \supset \psi \leq \neg\phi & \textit{or} \\ \neg\phi < \phi \supset \chi \quad \text{and} \quad \psi \leq \chi & \end{cases}$$

This condition applies only in the ordinary case when $\nvdash \neg\phi$. In the limiting case when ϕ is contradictory, the relativization of $\leq$ to $K * \phi = \mathcal{L}$ results in an unchanged entrenchment relation: $\psi \leq^{*C}_{\phi} \chi$ iff $\psi \leq \chi$.[13]

For the purposes of comparison with other approaches in Section 8, it is useful to take down that (CER) can equivalently be written thus:

$$\text{(CER}'\text{)} \qquad \psi \leq^{*C}_{\phi} \chi \;\text{ iff }\; \begin{cases} \phi \supset (\psi \wedge \chi) \leq \neg\phi \;\text{ and }\; \phi \supset \psi \leq \neg\phi & \textit{or} \\ \neg\phi < \phi \supset (\psi \wedge \chi) \;\text{ and }\; \psi \leq \chi & \end{cases}$$

Notice that the clauses $\phi \supset (\psi \wedge \chi) \leq \neg\phi$ and $\neg\phi < \phi \supset (\psi \wedge \chi)$ that determine a case distinction here simply mean $\psi \wedge \chi \notin K * \phi$ and $\psi \wedge \chi \in K * \phi$, respectively, in virtue of (Def $*$ from $\leq$). So even though (CER$'$) looks more complicated than (CER), it is perhaps easier to understand.

Can we capture the behaviour of such functions in terms of 'rationality postulates' for iterated belief revision? It turns out that we can.

6 A representation theorem

So far we have not considered any postulate that genuinely applies to iterated revision. Let us now add to the AGM axioms one additional postulate that is concerned with iterated revisions. We present it in a version that is close to the original AGM postulates, and a version that makes clear that the postulate is thought to apply not just to K, but to all $K * \Phi$ reachable from K through a sequence of revisions.

($*$9C) If $\neg\psi \in K * \phi$, then $(K * \phi) * \psi = K * \psi$.

($*$9C$^\omega$) If $\neg\chi \in K * (\Phi \cdot \psi)$, then $K * (\Phi \cdot \psi \cdot \chi) = K * (\Phi \cdot \chi)$.[14]

---

[13] If $K \dot{-} \phi$ is the $\leq$-based AGM contraction of K with respect to ϕ (Gärdenfors and Makinson 1988) and $K \ddot{-} \phi$ is the $\leq$-based severe withdrawal of K with respect to ϕ (Pagnucco and Rott 1999), and if $\leq^{\dot{-}}_{\phi}$ and $\leq^{\ddot{-}}_{\phi}$ are the relativizations of $\leq$ to $K \dot{-} \phi$ and $K \ddot{-} \phi$ respectively, then it can be shown that

$$\psi \leq^{\dot{-}}_{\phi} \chi \quad \text{iff} \quad \psi \leq^{*}_{\neg\phi} \chi \text{ or } \psi \notin K$$
$$\psi \leq^{\ddot{-}}_{\phi} \chi \quad \text{iff} \quad \psi \leq^{\dot{-}}_{\phi} \chi \text{ or } \psi \leq \phi$$

Since $\leq$ includes ties and incomparabilities, the relations standing on the left-hand sides are *coarsenings* of those standing on the right-hand sides.

[14] There is an alternative way of generalizing ($*$9C) to longer sequences, viz.: If $\neg\chi \in K * \psi$, then $K * (\psi \cdot \chi \cdot \Phi) = K * (\chi \cdot \Phi)$. For conservative belief change, no problem arises. Whether the prescription to work off the sequence "from left to right" gives the correct results, however, depends on the context in which it is used. See footnote 26.

For reasons that will become clear soon, we call a basic iterative revision function $*$ that satisfies ($*9C^\omega$) a *conservative iterative revision function.* The new condition makes sure that the revision function is *i-conservative*, where 'i' stands for '*i*nconsistent case' (input inconsistent with prior beliefs).

Postulate ($*9C^\omega$) may be viewed as a condition of *diachronic coherence* or *conservatism*. In the AGM model, the idea of minimal change is applied – so far as it is applied at all – only to pairs of prior and posterior theories. We extend this idea to whole sequences of theories, thereby introducing an obvious course-of-values dependency of revisions which is consonant with the idea of taking the doxastic history of an agent seriously (see part I of this paper). Like ($*3^\omega$) and ($*4^\omega$), the postulate ($*9C^\omega$) is orthogonal as it were to the conditions of dispositional coherence imposed by ($*7^\omega$) and ($*8^\omega$). Unlike ($*3^\omega$) and ($*4^\omega$), it addresses the case where the input is inconsistent with the prior belief set, and it therefore complements those postulates. Taken together with the AGM axioms ($*3^\omega$) and ($*4^\omega$), the new axiom ($*9C^\omega$) for iterated revisions in effect allows us to reduce iterated revision or belief sets to one-step belief revision. While ($*3^\omega$) and ($*4^\omega$) take care for the case where the input is consistent with the current beliefs (c-conservativity), ($*9C^\omega$) just addresses the case where it is inconsistent with them (i-conservativity). The three postulates taken together imply

$$(*3,4,9C) \quad K * \phi * \psi = \begin{cases} (K*\phi)+\psi & \text{if } \neg\psi \notin K*\phi \\ K*\psi & \text{otherwise} \end{cases}$$

The full version also describing the "prehistories" is

$$(*3,4,9C^\omega) \;\; K * (\Phi \cdot \psi \cdot \chi) = \begin{cases} (K*(\Phi\cdot\psi))+\chi & \text{if } \neg\chi \notin K*(\Phi\cdot\psi) \\ K*(\Phi\cdot\chi) & \text{otherwise} \end{cases}$$

The number of revision steps can be reduced one by one in this way, until we reach the point where the sequential revision is reduced to a revision by a single sentence. But notice that we cannot get rid of the operation of revision altogether.

Now we proceed to the main task of the present section and show that we have indeed found the postulate that characterizes the proposed belief revision strategies. The first theorem tells us that all basic iterated revision functions based on conservative revisions of an entrenchment relation satisfy ($*9C^\omega$).

Observation 8. (Soundness) Every iterated revision function $* : \mathcal{L}^\omega \to \mathbb{K}$ that is conservatively determined by a basic entrenchment relation $\leq$ satisfies ($*1^\omega$) – ($*6^\omega$) as well as ($*9C^\omega$).

Using Observation 6, part (d), and Observation 7, the following is an immediate consequence.

Corollary 9. Every iterated revision function $* : \mathcal{L}^\omega \rightarrow \mathbb{K}$ that is conservatively determined by a standard entrenchment relation $\leq$ satisfies ($*1^\omega$) – ($*9C^\omega$).

Finally we show that every iterated basic (or AGM) revision function that also satisfies ($*9C^\omega$) can be rationalized by an initial basic (standard) entrenchment relation and its conservative revisions.

Observation 10. (Completeness) For every iterated revision function $* : \mathcal{L}^\omega \rightarrow \mathbb{K}$ that satisfies ($*1^\omega$) – ($*6^\omega$) as well as ($*9C^\omega$), there is a basic entrenchment relation $\leq$ that conservatively determines $*$.

Using Observation 2, part (d), the following is an immediate consequence.

Corollary 11. For every iterated revision function $* : \mathcal{L}^\omega \rightarrow \mathbb{K}$ that satisfies ($*1^\omega$) – ($*8^\omega$) as well as ($*9C^\omega$), there is a standard entrenchment relation $\leq$ that conservatively determines $*$.

For the reasons mentioned in Section 4 of Part I of this paper, no use is made here of two-place revision functions. Such functions would save us the trouble of using sequences of input sentences as arguments since for two-place functions every resulting belief set could in the next revision step reappear as the first argument of the revision function. But it is indeed impossible to use such functions to simulate belief change based on conservative entrenchment revision. The reason lies in the importance of the "doxastic history" of the agent; one and the same belief set may be revised differently on different occasions, depending on its genealogy.

We consider the following example. We start from a state we call *complete ignorance*, i.e., from the unformed entrenchment relation that is defined by $\phi \leq \psi$ iff $\phi \notin Cn(\emptyset)$ or $\psi \in Cn(\emptyset)$, and from the ignorant belief set $K = \{\phi : \bot < \phi\} = Cn(\emptyset)$. Now consider $K * (p \cdot q)$ and $K * (q \cdot p)$. According to ($*3$) and ($*4$), the belief sets obtained are the same, namely $Cn(p, q)$. But the order of the revisions matters: The belief states are not identical! This can be seen from the behaviour of further revisions. Condition ($*3,4,9C$) tells us that

$$K * (p \cdot q \cdot \neg(p \wedge q)) = K * (p \cdot \neg(p \wedge q)) = Cn(p \wedge \neg q)$$

and

$$K * (q \cdot p \cdot \neg(p \wedge q)) = K * (q \cdot \neg(p \wedge q)) = Cn(\neg p \wedge q)$$

So the belief change potential of $K * (p \cdot q)$ and $K * (q \cdot p)$ are different here. Changes are sensitive to doxastic histories. We can also see from this example that the first piece of evidence is considered superior to the second piece of evidence, while the third one is strongest, due to ($*2$). The comparative recency of information does not translate systematically into comparative importance, strength or entrenchment. Let us brand this property as *temporal incoherence*. It will be discussed in a larger context in the next section.

As a related point, we lay stress on the fact that while there is basically a reduction of iterated belief revision to one-step belief revision as regards *belief sets* through ($*$3,4,9C$^\omega$), there is no such reduction as regards *belief states*.[15] For every sequence Φ there is a sentence ϕ such that $K * \Phi = K * \phi$, but it is in general not true that for a sequence Φ there is a sentence ϕ such that $*_\Phi = *_\phi$. For example, let us start from complete ignorance again. Then we have $K * (p \cdot q) = K * (p \wedge q)$ and $K * (p \cdot q \cdot \neg(p \wedge q)) = Cn\,(p \wedge \neg q)$, but $K * ((p \wedge q) \cdot \neg(p \wedge q)) = Cn\,(\neg(p \wedge q))$. The cumulative effects of repeated revisions on an entrenchment relation cannot be reproduced by a single one-step revision.

7 Boutilier's model and the criticism of Darwiche and Pearl

It turns out that something very close to conservative iterated revision functions has been studied earlier in a semantic setting by Craig Boutilier (1993, 1996). Boutilier represents a belief state by a connected ordering of models, something which exactly corresponds to a Grovean (1988) system of spheres. According to Boutilier's suggestion, a revision of such a belief state by ϕ is effected by taking the set of all minimal ("most plausible") ϕ-worlds, moving them to the very bottom of the ordering (to the very center of the system of spheres), declaring that they will now form the cluster of most plausible worlds (the innermost sphere), and leaving the rest of the ordering (the rest of the sphere system) completely unchanged. Boutilier discusses extensively the behaviour of this semantic construction and shows that postulate ($*$9C) is satisfied. He motivates his suggestion which he calls "natural revision" (1993) or "minimal conditional revision" (1996) by the idea that the changes to the ordering of worlds (models) as well as the changes to the set of conditionals validated by the states should be minimal. He also mentions briefly that this at the same time means a minimal change to plausibilities and entrenchments (1996, pp. 264, 273), but he does not elaborate on this remark. Besides the fact that Boutilier takes a semantic approach while we have based our considerations on the notion of epistemic entrenchment, there

[15] Boutilier (1993, 1996) very carefully studies the reduction of repeatedly revised belief sets, but he does not emphasize that the revisions of the belief states are not reducible in the way indicated. Boutilier (1993, Corollary 4, 1996, Corollary 8) replaces '$(K * \phi) + \psi$' in condition ($*$3,4,9C) by '$K * (\phi \wedge \psi)$' – which are identical only if one presupposes the validity of ($*$7) and ($*$8). As emphasized before, we want to study a form of diachronic coherence without presupposing any kind of dispositional coherence. It will be argued below that in the more general context, the two conditions actually stand for quite different ideologies.

are two essential differences. First, Boutilier does not formulate a completeness theorem. Darwiche and Pearl (1994, Theorem 1; 1997, Theorem 11) supply the missing completeness theorem for Boutilier's revision operator. Second, and more importantly, Boutilier works in a framework that always presumes full comparability of the models (with respect to his ordering), and therefore he is committed to accepting the AGM postulates ($*7$) and ($*8$). Similarly, Darwiche and Pearl's representation theorem is proved in the special context where the full power of AGM revision functions is available.[16] By contrast, it is one of the major aims of the present paper to show that the issue of *diachronic coherence* in iterated belief revision can be completely separated from the characteristic postulates ($*7$) and ($*8$) of the AGM theory which embody *dispositional coherence*. The relevant observations can all be made even in a context in which nothing about dispositional coherence is presupposed.[17]

By way of criticism, Darwiche and Pearl came up with a critical example. If an agent first believes that some observed animal is a bird, then sees that the animal is red, but is finally told that the animal is not a bird after all, then ($*9$C) prescribes that he must give up his belief that the animal is red. More formally, if we assume that ψ is consistent with $K * \phi$, then conservative belief change yields

$$K * (\phi \cdot \psi \cdot \neg\phi) = K * \neg\phi$$

even if intuitively ψ has "nothing to do" with ϕ.[18] In my view, this little story comes very close to a falsification of postulate ($*9$C). Boutilier's (1996, p. 296) reply that minimal conditional revision "can still be viewed as a computationally attractive approximation method for ... 'ideal' revision" sounds evasive and comes close to an admission of defeat.

So it seems a good stategy to relax the conservative constraints. Darwiche and Pearl (1994, 1997) proposed to give up ($*9$C) in favour of four postulates that taken together are substantially weaker than the former. It has been remarked several times that even the Darwiche-Pearl postulates are incompatible with the

[16] In claiming the equivalence of Boutilier's model with conservative entrenchment revision including ($*7$) and ($*8$), I rely on Darwiche and Pearl's representation result and the one presented in this paper. It does not seem necessary here to prove the equivalence by means of the usual bridge constructions connecting systems of spheres and epistemic entrenchment relations. Such bridge principles have essentially been known since Grove (1988) and Gärdenfors (1988); for a recent exposition of this problem, see Pagnucco and Rott (1999).

[17] Compare Part I of this paper. It is important here to see that the basic postulates plus those of c-conservativity have nothing to do with dispositional coherence, notwithstanding the fact that AGM have shown that these postulates entail, e.g., the existence of a choice function generating such revisions. The point is that this choice function, and thus the agent's dispositions, need not be coherent in the sense of the theory of rational choice.

[18] The order of the revision steps is important. If we reverse the first two revisions, then $K * \psi * \phi * \neg\phi = (K * \psi) + \neg\phi$ which is a plausible result.

AGM postulates $(*1)$ through $(*8)$.[19] This is a striking observation; in the proof of the corollary to Theorem 10 we have given an explicit *construction* that satisfies all the AGM postulates plus $(*9\text{C})$ which is—given the AGM postulates—stronger than the Darwiche-Pearl postulates. Therefore that set must be consistent! How can this be? The solution to this little logical puzzle is that Darwiche and Pearl, as well as the other authors mentioning the incompatibility, take all the AGM postulates to be postulates for two-place revision functions that take as arguments pairs consisting of a belief set and a sentence. This immediately entails the Ahistoricity which is blatantly violated by the model suggested in this paper (as it should be, I think). I have argued at length in Part I of this paper that the approach to iterated belief change using two-place functions is inappropriate. In our setting with unary revision functions taking sequences, the Darwiche-Pearl postulates become

DP1$^\omega$	If $\psi \in Cn(\chi)$, then $K * (\Phi \cdot \psi \cdot \chi) = K * (\Phi \cdot \chi)$
DP2$^\omega$	If $\neg\psi \in Cn(\chi)$, then $K * (\Phi \cdot \psi \cdot \chi) = K * (\Phi \cdot \chi)$
DP3$^\omega$	If $\psi \in K * (\Phi \cdot \chi)$, then $\psi \in K * (\Phi \cdot \psi \cdot \chi)$
DP4$^\omega$	If $\neg\psi \notin K * (\Phi \cdot \chi)$, then $\neg\psi \notin K * (\Phi \cdot \psi \cdot \chi)$

Notice a difference in the antecedents: In $(*9\text{C}^\omega)$, we have the precondition that $\neg\chi$ is in $K * (\Phi \cdot \psi)$, while in the related condition (DP2$^\omega$) the precondition is equivalent to requiring that $\neg\chi$ is in $Cn(\psi)$. Since $Cn(\psi)$ is a subset of $K*(\Phi\cdot\psi)$ (by $(*1^\omega)$ and $(*2^\omega)$), and since $Cn(\psi)$ may in general be expected to be a rather small subset of $K*(\Phi\cdot\psi)$, the antecedent of $(*9\text{C})$ will be much more frequently satisfied than the antecedent of (DP2$^\omega$) – a fact that sheds some light on the strength of $(*9\text{C})$.

Since in our framework postulates for belief revision do *not* quantify over all belief sets K, there is no incompatibility between the AGM and the Darwiche-Pearl postulates. One might expect that conservative belief revision functions in our sense, i.e., functions satisfying postulate $(*9\text{C})$, *a fortiori* satisfy all of the Darwiche-Pearl postulates. Since, however, it is part of the philosophy of this paper to renounce a commitment to dispositional coherence, i.e., to $(*7)$ and $(*8)$, we cannot simply appeal to the results of Darwiche and Pearl, but must conduct an analysis of our own. So we focus on basic revision functions that are c-conservative and i-conservative. It turns out that dispositional coherence is not needed for (DP2$^\omega$) – (DP4$^\omega$), but that it is necessary for (DP1$^\omega$).

Observation 12. (a) Every iterated revision function $* : \mathcal{L}^\omega \to \mathbb{K}$ that satisfies $(*1^\omega)$ – $(*6^\omega)$ as well as $(*9\text{C}^\omega)$ also satisfies (DP2$^\omega$) – (DP4$^\omega$);

[19] Freund and Lehmann (1994, p. 8), Nayak et al. (1996, p. 121), and Darwiche and Pearl (1997, p. 7).

(b) every iterated revision function $* : \mathcal{L}^\omega \to \mathbb{K}$ that satisfies $(*1^\omega)$ – $(*9C^\omega)$ in addition satisfies $(DP1^\omega)$.

Since part (b) cannot be strengthened so as to apply in the context of revisions without dispositional coherence, Darwiche and Pearl do not propose simply a weakening of conservative belief change form the point of view taken in this paper. They also strengthen it in one respect. Previous discussions have suggested that (DP2) is problematic because of its alleged incompatibility with the AGM postulates. From our perspective, (DP1) is more questionable than (DP2). But Darwiche and Pearl's argument that $(DP1^\omega)$ – $(DP4^\omega)$ do not imply $(*9C)$ of course remains valid, and so does their criticism of the latter condition.

Apart from the particular type of example given by Darwiche and Pearl, one can advance a more general criticism of the strategy involved in iterated belief revision obeying $(*9C)$. If ϕ is a new piece of information in a belief revision process, it has a truly privileged status: It is to be accepted, by force of the condition of Success, $(*2)$. Once accepted, however, ϕ has has a very weak doxastic status.[20] It is an immediate consequence of K-minimality and the definition of conservative entrenchment revision that if ϕ is a "new" belief that had not been in in the prior belief set $K = K_\leq$ to begin with, then $\phi \leq^{*C}_\phi \psi$, for all ψ in the posterior belief set $K * \phi = \{\xi : \neg\phi < \phi \supset \xi\}$. So although we invariably accept a new piece information, we do so only at the lowest level of entrenchment possible. Due to K-representation, ϕ is actually *less* entrenched than all the "old" beliefs $\psi \in K$ that have been retained in the revised belief set $K * \phi$.

This has severe consequences when more information, ψ say, comes in. Condition $(*9C)$ tells us that previous revisions are undone when the agent works off a sequence of inputs backwards step by step, as long as ψ is found to be in conflict with the results of these earlier revisions. So while new information gets the highest priority possible, it is given up all too readily when more information comes in. Such a strategy of belief change is open to criticism on two counts. First, it shows no consistent attitude towards the novelty of information. The most recent input sentence is always embraced without reservation, the last but one input sentence, however, is treated with utter disrespect. Second, even if we had a coherent strategy of letting entrenchments vary in a principled way with the time at which a piece of information was received, it is not clear how this covariance might be justified.

The shortcomings of the conservative recipe notwithstanding, it does have its advantages. It shows that the argument given at the beginning of Section 4 is flawed: One *can* use the old entrenchment relation for the beliefs in the revised belief set. The method proposed shows that it is possible to formalize a notion of conservatism that transcends the usual sticking to one's old beliefs. And

[20]Spohn (1988, p. 114) made the same observation in a context with dispositional coherence.

the method is in an instructive way different both from (i) taking one and the same belief-independent entrenchment relation for all belief sets and from (ii) 'substantially' revising the inner structure of the entrenchment relation.[21]

Still Darwiche and Pearl have advanced good reasons for thinking that we should not be as conservative as we have been up to now. What to do in this situation? We shall now survey three principal alternatives to conservative belief change in a non-numerical framework: external revision, radical revision and moderate revision.

8 Comparison with other models for non-quantitative iterated belief revision

It was only with a delay of about a decade that belief revision theory faced the challenge that is raised by iterations of belief change. Not the AGM trio themselves, but other researchers have shaped the landscape here.[22] Among the people who have made significant contributions to the development of the field that I am going to survey are W. Spohn and S.O. Hansson, C. Boutilier, M. Freund and D. Lehmann, A. Nayak and his coauthors, A. Darwiche and J. Pearl, N. Friedman and J. Halpern, K. Segerberg and J. Cantwell, O. Papini, S. Konieczny and R. Pino Pérez, C. Areces and V. Becher (I do not claim completeness for this list). The discussion in this section is not meant to aim at a comprehensive and fully explicit exposition of alternative approaches (we will have to skip all the proofs).

Many researchers have suggested to use richer representations of doxastic states than just belief sets. For instance, Spohn (1988) revises ordinal conditional functions (now often called "ranking functions" or "κ-rankings"), Nayak (Nayak 1994, Nayak, Nelson and Polansky 1996, Nayak *et al.* 1996) revises entrenchment relations,[23] Boutilier (1993, 1996) revises "revision models", Darwiche and Pearl (1994, 1997) revise primitive "epistemic states", Cantwell (1999), following the lead of Segerberg, revises "hypertheories". My aim in this paper has been to stay closer to the original AGM account in this paper. Although I view a doxastic state as a unary iterated revision function, my proposal can also be regarded as one for potentially revising a fixed belief set $K = *(\mathbf{0})$ step by step, by finite

[21] In the long run, after many items of input have been processed, the entrenchment relation gets substantially revised by conservative revision as well. We will discuss methods of substantially revising entrenchments in one step presently.

[22] As far as I am aware, the only thing that AGM have written about iterated belief change is one section in Alchourrón and Makinson's (1985) paper on safe contraction.

[23] Like in Nayak's papers, entrenchment relations are taken to be epistemic states in Rott (1991b) where entrenchment relations are revised in a derivative way through changes of conditional knowledge bases.

sequences of sentences. In AGM-style modellings, mental attitudes that underlie belief revision behaviour are not represented in the object language, but in the metatheory. In any case, we somewhere need full-blown formal representations of doxastic states, including structures (such as entrenchment relations) that encode an agent's dispositions to change his doxastic states.

8.1 External belief revision

One way of dealing with the intricacies of iterated belief change is to suggest essentially taking a God's eye point of view in which the revision-guiding structures are *external* to the agent's belief state. (This is approach (i) just mentioned.) Some authors have presented a fixed "hierarchy" (Alchourrón and Makinson 1985), a fixed "epistemic preference relation" (Schlechta 1991) or a fixed "generalized epistemic entrenchment relation " (Rott 1992) which is independent of the agent's present belief set, and suggested to use it for varying belief sets. The approach has been advocated in a more principled way by Hansson (1992, 1999), Freund and Lehmann (1994) and in particular Areces and Becher (2001).

We need to assume that we have a fixed entrenchment relation $\leq_F$ for a fixed theory $F = \{\phi : \bot < \phi\}$ which is independent of the agent's current belief set. F may be God's theory (the one true $\mathcal{L}$-theory about the real world), but it may also be any other theory, even the inconsistent theory will do perfectly well.[24] Given this idea, the relevant prescription for the revision of epistemic entrenchment, usable for iterated belief change starting from a belief set K, is easily specified.

(EER)
$$\psi \leq^{*E}_{\phi} \chi \text{ iff } \begin{cases} \psi \notin K * \phi & \textit{or} \\ \chi \in K * \phi \text{ and } \psi \leq_F \chi \end{cases}$$
$$\text{iff } \begin{cases} \phi \supset \psi \leq \neg\phi & \textit{or} \\ \neg\phi < \phi \supset \chi \text{ and } \psi \leq_F \chi \end{cases}$$

The entrenchment relation used, $\leq_F$, does not depend on either the old belief set K or the new belief set $K * \phi$. Notice that this is the only formal difference with conservative entrenchment revision. The second 'iff' applies only when $\not\vdash \neg\phi$. In the limiting case when ϕ is contradictory, the relativization of $\leq$ to $K * \phi = \mathcal{L}$ results in the external entrenchment relation: $\psi \leq^{*E}_{\phi} \chi$ iff $\psi \leq_F \chi$.

The corresponding postulate for external iterated belief change that is captured by this entrenchment operation is

(*9E) If $\neg\phi \in K$, then $K * \phi = F * \phi$.

(*9E$^\omega$) If $\neg\chi \in K * (\Phi \cdot \psi)$, then $K * (\Phi \cdot \psi \cdot \chi) = F * \chi$.

[24] Use of the inconsistent theory is in fact suggested by Rott (1992) and Areces and Becher (2001).

Except for the unproblematic case where c-consistency applies, $*$ does not depend on the current belief set K. So we really have postulates here that are applicable to iterated belief change. The orderings or other devices for belief change that are not related to the agent't belief set do not really encapsulate *the agent's* preferences or entrenchments. For instance, in such a model some of the agent's beliefs will in general be less "preferred" or "entrenched" than some of his non-beliefs. This seems to be a somewhat abusive application of the names of these relations. But more importantly, as ($*9\mathrm{E}^{\omega}$) brings out clearly, such a model defines a two-place revision function, and as such it shows no sensitivity at all to the history of belief changes. To my mind, external revision embodies a bad philosophy. First, it is not at all clear where the external relation $\leq_F$ comes from and how it is to be interpreted. And second, the history of belief changes should have some potential effect in the dynamics of belief.

We will now address two methods that substantially revise the inner structure of entrenchment relations (cf. method (ii) in Section 7).

8.2 Radical belief revision

Taking up a point from the end of Section 7, we recall that in conservative belief revision, acceptance of new information is forced, but it is acceptance with minimal commitment. In a sense, therefore, conservative entrenchment revision can be regarded as a dual to the following idea:

$$\psi \leq^{*R}_{\phi} \chi \text{ iff } \phi \supset \psi \leq \phi \supset \chi$$

While conservative entrenchment revision assigns the input sentence ϕ minimal entrenchment among the beliefs, this suggestion makes ϕ maximally entrenched, and in fact irremediably so: Clearly, $\psi \leq^{*R}_{\phi} \phi$ for all ψ, and this even remains so for all subsequent revisions of $\leq^{*R}_{\phi}$. Let us call this method *radical entrenchment revision*. By Extensionality, radical belief change preserves the prior entrenchment relation within the set $Cn(\{\neg\phi\})$, but it extinguishes all differences of entrenchment within the set $Cn(\{\phi\})$. Clearly this method violates the Maximality condition for entrenchment relations, and correspondingly, the consistency condition ($*5$) for revisions gets violated as well. Radical belief change was first ventilated (but ultimately rejected) in Rott (1991a), and has recently been endowed with an axiomatization by Fermé (2000). The method has found supporters who argue that it represents the right way to go about iterated belief changes in two particular contexts: in *hypothetical belief change* (Segerberg 1998) and in belief change occasioned by *observations* that are taken as irrefutable knowledge (Friedman and Halpern 1996, 1999). An important point that is emphasized by these authors (for different reasons) is that the sentences in a sequence of hypotheses or observations must not or cannot be inconsistent with each other.

The characteristic postulate of the radical method of belief revision is simply

$$K * \phi * \psi = K * (\phi \wedge \psi)$$

A truly general method of iterated belief change should not systematically violate the crucial consistency condition ($*5$). So the above definitions are too radical. Slightly less radical are the following, official definitions of revised entrenchment relations

(RER) $\quad \psi \leq_{\phi}^{*R} \chi \quad \text{iff} \quad \phi \supset \psi \leq \phi \supset \chi \;\text{ and }\; \text{either} \;\; \not\vdash \psi \text{ or } \vdash \chi$

and or iterated revision operations

(*9R) $$K * \phi * \psi = \begin{cases} K * (\phi \wedge \psi) & \text{if } \neg\psi \notin Cn(\phi) \\ Cn(\psi) & \text{otherwise} \end{cases}$$

(*9R$^{\omega}$) $$K * (\Phi \cdot \psi \cdot \chi) = \begin{cases} K * (\Phi \cdot (\psi \wedge \chi)) & \text{if } \neg\chi \notin Cn(\psi) \\ Cn(\chi) & \text{otherwise} \end{cases}$$

The method characterized by these conditions is still fairly radical. In $\leq_{\phi}^{*R}$, all distinctions between propositions inconsistent with ϕ are wiped out. Radical belief change is another limiting case that is certainly not better than the temporally incoherent method of conservative belief change. We still have to ask: What to do in this situation?

8.3 Moderate belief revision

Obviously, we need to steer a middle course between the extremes of conservatism and radicalism. Fortunately, there is one. I call it – because it is "between" the conservative and the radical solution, "moderate revision". I have taken inspiration from the work of Nayak (1994) here, but the method seems somehow to have made its way into the folklore whithout there being a classical reference or an established name for this method. It was probably first mentioned, but rejected by Spohn (1988, pp. 112–113).[25] After Nayak who seems to have been the first to advocate and study the method, it (or a close variant of it) has been mentioned or rediscovered more recently by quite a number of authors: Liberatore (1997, p. 280) calls it "prioritized iteration of revision", Glaister (1998, pp. 30–34) calls it "J-revision" ('J' for 'Jeffrey'), Kelly (1999, p. 16) calls (a numerical variant of) this method "the lexicographic operator", Konieczny (1998, pp. 7–8) and

[25] Spohn rejected the method because (a) it is not reversible, (b) it is not commutative and (c) it always gives primacy to the last input. I doubt whether one should consider reversibility and commutativity as necessary conditions for reasonable revision operations. I agree that the last input should not always be given primacy, but it seems to me that there can and indeed should be *some* revision operation that shows this behaviour.

Konieczny and Pino Pérez (2000, pp. 353–355) call it "basic memory operator", and Papini (2001, pp. 285–289) uses the symbol '$\circ_{\triangleright}$' without introducing any special name.

Nayak's (1994) model works with connected and transitive entrenchment relations and thus subscribes to dispositional coherence. An important generalization that he studies is the case where the input is not in form of a sentence but in the form of another entrenchment relation. Interesting though it is, we do not need this additional generality for the task before us. We only need Nayak's special case of "naked evidence", in which an input sentence ϕ is identified with the entrenchment relation $\leq^{\phi}$ that puts $\psi \leq^{\phi} \chi$ iff either $\vdash \chi$, or $\phi \nvdash \psi$, or both $\phi \vdash \chi$ and $\nvdash \psi$. Simplified and written out explicitly, Nayak's (1994, p. 372) method of revising a given entrenchment relation $\leq$ by naked evidence ϕ (that is, by $\leq^{\phi}$) is equivalent to this condition:

$$\text{(MER)} \qquad \psi \leq^{*M}_{\phi} \chi \text{ iff } \begin{cases} \psi \wedge \chi \notin Cn(\phi) \text{ and } \phi \supset \psi \leq \phi \supset \chi & \textit{or} \\ \psi \wedge \chi \in Cn(\phi) \text{ and } \psi \leq \chi \end{cases}$$

Nayak's suggestion is a compromise between conservative belief change and radical belief change. The compromise is even recognizable in the definition. In the first line of Nayak's case distinction we can recognize the characteristic clause of radical belief change, in the second line we recognize the characteristic clause of conservative belief change. Radical entrenchment change is as it were restricted from universal applicability to application outside $Cn(\phi)$, whereas the characteristically conservative reutilization of the old entrenchment relation is restricted from $K * \phi$ to $Cn(\phi)$ (which is a subset of the former, by (∗1) and (∗2)). Let us call this method *moderate entrenchment revision*. $\leq^{*M}_{\phi}$ is less conservative than conservative entrenchment revision $\leq^{*C}_{\phi}$ in that it re-uses $\leq$ not in the whole of $K * \phi$, but only in $Cn(\phi)$ which is a subset of the former. It is more conservative than radical entrenchment revision $\leq^{*R}_{\phi}$ in that it preserves $\leq$ not only within $Cn(\neg\phi)$, but also within $Cn(\phi)$.

Entrenchment relations bequeath their properties to their moderately revised descendants:

Observation 13. If $\leq$ satisfies Reflexivity (or respectively, Extensionality, Choice, Maximality), so does the moderately revised relation $\leq^{*M}_{\phi}$. If $\leq$ is a transitive basic entrenchment relation, so is $\leq^{*M}_{\phi}$.

Like the condition for entrenchments, the condition for iterated revision functions mirrors the fact that moderate belief change takes a middle course between conservative and radical belief change. The characteristic postulate of Nayak's moderate method is

$$(*9\mathrm{M}) \qquad K * \phi * \psi = \begin{cases} K * (\phi \wedge \psi) & \text{if } \neg\psi \notin Cn(\phi) \\ K * \psi & \text{otherwise} \end{cases}$$

$$(*9\mathrm{M}^\omega) \qquad K * (\Phi \cdot \psi \cdot \chi) = \begin{cases} K * (\Phi \cdot (\psi \wedge \chi)) & \text{if } \neg\chi \in Cn(\psi) \\ K * (\Phi \cdot \chi) & \text{otherwise}^{26} \end{cases}$$

In accordance with the idea of the success postulate, $(*2)$, moderate belief change gives priority to incoming information. Unlike conservative entrenchment revision, however, it keeps on paying a lot of respect to the recency of information in subsequent revisions. It is easy to check that it comes to terms with Darwiche and Pearl's red bird example (the animal is believed to be red even if it turns out that it is not a bird). Moderate belief change embodies a principled attitude towards the doxastic value of novelty in a sequence of inputs: "The more recent a piece of information is, the better."[27] For this reason, it seems to me clearly better than conservative belief change, and indeed the best of all the iterated strategies that we have discussed in this paper.

Now we show that the postulate $(*9\mathrm{M}^\omega)$ characterizes the moderate belief revision strategy. The first theorem tells us that all basic iterated revision functions based on moderate entrenchment revisions satisfy $(*9\mathrm{M}^\omega)$.

Observation 14. (Soundness) Every iterated revision function $* : \mathcal{L}^\omega \rightarrow \mathbb{K}$ that is moderately determined by a basic entrenchment relation $\leq$ satisfies $(*1^\omega)$ – $(*2^\omega)$ and $(*5^\omega)$ – $(*6^\omega)$ as well as $(*9\mathrm{M}^\omega)$. However, it does not generally satisfy $(*3^\omega)$ and $(*4^\omega)$, nor does it generally satisfy $(*7^\omega)$ and$(*8^\omega)$

Using Observation 6, part (d), and Observation 7, the following is an immediate consequence.

Corollary 15. Every iterated revision function $* : \mathcal{L}^\omega \rightarrow \mathbb{K}$ that is moderately determined by a standard entrenchment relation $\leq$ satisfies $(*1^\omega)$ – $(*9\mathrm{M}^\omega)$.

Conversely, we show that every iterated basic (or AGM) revision function that also satisfies $(*9\mathrm{M}^\omega)$ can be rationalized by an initial basic (standard) entrenchment relation and its moderate revisions.

---

[26]There is an alternative way of generalizing the lower line of $(*9\mathrm{M})$ to longer sequences, viz.: If $\neg\chi \in K * \psi$, then $K * (\psi \cdot \chi \cdot \Phi) = K * (\chi \cdot \Phi)$. In the context of moderate belief change, then, the results are different (and indeed not as intended), as the following example shows: Let ϕ, ψ and χ three sentences that are pairwise consistent but jointly inconsistent. Then in moderate belief change we get $K * \phi * \psi * \chi = K * \phi * (\psi \wedge \chi) = K * (\psi \wedge \chi)$, while the from-left-to-right alternative condition would give us $K * \phi * \psi * \chi = K * (\phi \wedge \psi) * \chi = K * \chi$. Compare footnote 14.

[27]Papini (2001) also studies the dual operation where older information is always held in higher regard than novel information.

Observation 16. (Completeness) For every iterated revision function $* : \mathcal{L}^\omega \to \mathbb{K}$ that satisfies $(*1^\omega)$ – $(*2^\omega)$ and $(*5^\omega)$ – $(*6^\omega)$ as well as $(*9\mathrm{M}^\omega)$, there is a basic entrenchment relation $\leq$ that moderately determines $*$, except that $K = *(\mathbf{0})$ need not be identical with $K_\leq$.

Using Observation 2, part (d), the following is an immediate consequence.

Corollary 17. For every iterated revision function $* : \mathcal{L}^\omega \to \mathbb{K}$ that satisfies $(*1^\omega)$ – $(*8^\omega)$ as well as $(*9\mathrm{M}^\omega)$, there is a standard entrenchment relation $\leq$ that moderately determines $*$.

Conservative and moderate belief revision represent very different "ideologies". Let us look at three aspects of this difference.

First, it is important to point out that in general $K * (\phi \wedge \psi)$ is different from $(K * \phi) + \psi$. Of course, in the AGM context with dispositional coherence (i.e., $(*7)$ and $(*8)$), the two come out as the same, as long as ψ is consistent with $K * \phi$. But they give different results indeed if we cannot presume dispositional coherence.

Second, even it is straightforward to perform the operation of forming an expansion '+' on sets of beliefs, we have not yet provided an account of what it means to perform an *expansion* of an epistemic state. On the belief set level, everything is as trivial as expansions should be. The operation '+' means set-theoretic addition of the new piece of information, followed by logical closure; in the belief-contravening case this leads to an inconsistency. For the iterated case we get

$$(+9\mathrm{C}) = (+9\mathrm{M}) \qquad K + \phi + \psi = K + (\phi \wedge \psi)$$

However, the notion of an expansion for belief-revision guiding structures is not determined by this. Fortunately, it is rather straightforward to transform the ideas underlying (CER) and (MER) to the case of expansion which lead to inconsistent theories in the case of a belief-contravening information. Here is *conservative expansion*, expressed both in terms of one-step entrenchment revisions and in terms of iterated belief set revisions:

$$(\mathrm{CEE}) \qquad \psi \leq^{+C}_\phi \chi \ \text{ iff } \begin{cases} \phi \supset \psi \leq \bot & \textit{or} \\ \bot < \phi \supset \chi \ \text{ and } \ \psi \leq \chi \end{cases}$$

Alternative formulation:

$$(\mathrm{CEE}') \qquad \psi \leq^{+C}_\phi \chi \ \text{ iff } \begin{cases} \psi \wedge \chi \notin K + \phi \ \text{ and } \ \phi \supset \psi \leq \bot & \textit{or} \\ \psi \wedge \chi \in K + \phi \ \text{ and } \ \psi \leq \chi \end{cases}$$

The differences with (CER) are easy to detect. Where we now have '$\bot$' and '$+$', we used to have '$\neg\phi$' and '$*$'. Both (CEE) and (CEE$'$) show clearly that if ϕ is inconsistent with K (i.e., if $\bot < \neg\phi$), then the entrenchment relation does not change at all, that is, $\leq_\phi^{+C} = \leq$. The only difference in this case is that instead of $K = \{\psi : \bot < \psi\}$, it is the set $\mathcal{L}$ which is the belief set "associated with" $\leq_\phi^{+M}$.

Now let us turne to *moderate expansion*, expressed both in terms of one-step entrenchment revisions and in terms of iterated belief set revisions:

$$\text{(MEE)} \qquad \psi \leq_\phi^{+M} \chi \;\text{ iff }\; \begin{cases} \psi \wedge \chi \notin Cn(\phi) \;\text{ and }\; \phi \supset \psi \leq \phi \supset \chi & \textit{or} \\ \psi \wedge \chi \in Cn(\phi) \;\text{ and }\; \psi \leq \chi \end{cases}$$

The interesting point here is that there is no difference whatsoever between (MEE) and (MER), or between $\leq^{+M}$ and $\leq^{*M}$. The only difference is that if ϕ is inconsistent with K (i.e., if $\bot < \neg\phi$), then instead of $\{\psi : \bot <_\phi^{+M} \psi\}$, it is the set $\mathcal{L}$ which is the belief set "associated with" $\leq_\phi^{+M}$.

It is worth pointing out that the possibility of conservative and moderate expansion, as well as the subtleties of their differences, show that the revision of inconsistent theories does not present special difficulties for belief revision theories that work with closed belief sets. This topic cannot serve as a decisive argument in favour of theories that draw on the syntactical structure of belief bases.[28] In the evolution of belief sets that generated by conservative and moderate change operations (both expansions and revisions), inconsistent theories are just passing stages that are easily left in a subsequent revision.

Third, there is a difference that prima facie seems to be of little importance, but that does in fact represent a fundamental difference in the possible architectures of iterated belief revision processes. Let us compare ($*$9C) with ($*$9M). In conservative belief change, we have to check whether ψ is consistent with $K * \phi$, while in moderate belief change, we have to check whether ψ is consistent with ϕ. In a way, one can say that conservative belief change is *left-associative* because it first combines K with ϕ into a new belief set, and only in a subsequent step is this new belief set is revised by ψ, and so on: $K*\phi*\psi = (K*\phi)*\psi$. Straightforward as this is, moderate belief change indicates that it is not the only way to go. In ($*$9M$^\omega$), the antecedent check only relates to the two pieces of evidence, ϕ and ψ, and sees whether they can be conjoined or whether the latter overrules the former. In this

[28]Compare Hansson (1999, Chapter 3.4) who points out that certain methods that make good sense for belief bases become useless if applied to belief sets. Note that the meaning of Hansson's term "external revision" is different the meaning that we have attached to the term in this paper. Hansson's (1999, Chapter 3.2) proposal to employ *global* belief change functions comes down to what we have criticized as ahistoric (two-place) belief change functions in Parts I and II of this paper, and I my view the problems highlighted by Hansson should be blamed on this proposal rather than on the use of belief sets.

sense, moderate belief change is *right-associative*: $K * \phi * \psi = K * (\phi \cdot \psi)$.[29]

Philosophically, the moderate option seems better motivated than the conservative one, since in $K * \phi * \psi$ we have one prior set of sentences (which may include "soft" expectations and prejudices) and two pieces of input (which are supposed to be "hard", as reflected by the success condition $(*2)$). Seen from this angle, ϕ and ψ "belong together" more closely than K and ϕ Thus it appears that $K * \phi * \psi$ should be bracketed like $K * (\phi \cdot \psi)$ rather than the more commonly used $(K * \phi) * \psi$. And moderate belief change according to $(*9\mathrm{M}^{\omega})$ combines the first input χ with the second input ψ, while in conservative belief change according to $(*9\mathrm{C}^{\omega})$ the first input ψ is absorbed by the belief set K and is later, in $K * \psi$, not remembered as a piece of evidence that is worth special epistemic treatment.

But we have to be careful not to make too much of this difference! It is true that in condition $(*9\mathrm{M}^{\omega})$, we "work from right to left" in order to get the right results. So, if we look at the level of *belief sets*, the method is indeed right-associative. But in the manipulation of *entrenchments* or other belief-revision guiding structures (like systems of spheres or choice functions sketched in the next section), the operation is perfectly left-associative, as a brief reflection on (CEM) confirms.

8.4 Possible-worlds semantics for iterated belief change: Systems of spheres and semantic choice functions

We have mentioned earlier that there is a semantics of entrenchments in terms of nested systems of spheres (Grove, following David Lewis) or equivalently, connected orderings of possible worlds (Boutilier). The representation in terms of systems of spheres \$ which is the one which is most easily visualized. It must be pointed out, however, that the prize to be paid for ease of visualization is that systems of sphere are only appropriate in contexts where dispositional coherence, i.e., postulates $(*7)$ and $(*8)$, can be taken for granted.

Since we have aimed at factoring out dispositional coherence as encoded in $(*7)$ and $(*8)$, however, we cannot presume the existence of a (well-behaved) preference relation. But it is possible to work with *choice functions* over elementary sets of models or "possible worlds." Such a semantic choice function γ selects, for every set of possible worlds, the subset of those which are considered most plausible.[30]

[29] Another way of expressing the same thing is to say that conservative belief revision proceeds in an *integrative* or *stepwise* manner (every single piece of input occasions a transition from the prior to some posterior belief state), while moderate belief change follows an *evidential* or *sequential* strategy (the sequence of inputs can be processed before actually performing the revision operation).

[30] For details on this notion, see Rott (1998, 2001).

The last two colums of Table 1 summarize how choice functions and systems of spheres get revised according to the four methods of iterated belief change that we have been considering. In the table, $[\![\phi]\!]$ and $[\![K]\!]$ denote the sets of models or worlds verifying ϕ and all elements of K, respectively. In the presentation of external revision, γ_F and $\$_F$ denote the external choice function and the external system of spheres that are independent of the agent's belief set.

Since we have been working with epistemic entrenchment for the most part of this paper, it may be useful to give the semantics underlying entrenchment relations.[31]

In the more general account using choice functions, $\phi \leq \psi$ means that either $[\![\neg\phi]\!] \cap \gamma([\![\neg(\phi \wedge \psi)]\!]) \neq \emptyset$ (the principal case) or $\gamma([\![\neg(\phi \wedge \psi)]\!]) = \emptyset$ (the limiting case). Intuitively, this says that ϕ is not more entrenched than ψ if and only if among the set of most plausible worlds violating at least one of ϕ and ψ there are some worlds violating ϕ, provided there are any such worlds.

In the more demanding model using systems of spheres, $\phi \leq \psi$ means that all spheres of possible worlds that are completely covered by $[\![\phi]\!]$ are also covered by $[\![\psi]\!]$, or, essentially equivalently, that the smallest sphere intersecting $[\![\neg\phi]\!]$ is a subset of the smallest sphere intersecting $[\![\neg\psi]\!]$. This roughly says that the agent has to move "farther away" from what he actually believes in order to accommodate ψ than in order to accommodate ϕ.

Both of these accounts serve not only to form a better conception of entrenchment, but also make clear that the central condition of (Choice) for entrenchment which looks a bit weird at first sight is actually very natural.

9 Conclusion

In part I of this paper, conservatism was characterized as a type of diachronic cohrence that can be distinguished from other kinds of corehence in belief change, namely synchronic coherence (logical consistency and closure) and dispositional coherence (coherence across different potential revisions).

It became clear the conservatism plays a much more modest role in traditional belief revision theories than a wide-spread picture has suggested. There are only conditions that pertain to conservatism for changes that are not belief-contravening, viz. conditions (∗3) and (∗4). In this paper, we have aimed at strengthening AGM's 'c-conservatism' to a conservatism that applies to the case of belief-contravening changes. Such changes, after all, are the ones that make belief revision theories interesting.

It turned out that in order to deal with this case, it is not sufficient to follow the slogan '*Preserve as many beliefs as possible.*' Instead, we had to embrace the

[31] Compare Pagnucco and Rott (1999) and Rott (2001, Section 8.6).

	Iterated revision $K * \phi * \psi =$	Entrenchment $\psi \leq^*_\phi \chi$ iff	Semantic choice $\gamma^*_\phi(\llbracket\psi\rrbracket) =$	System of spheres $\$^*_\phi =$
Model requires	(∗1),(∗2),(∗5),(∗6)	(∗1),(∗2),(∗5),(∗6)	(∗1),(∗2),(∗5),(∗6)	(∗1)–(∗8), full comparability
External	$(K * \phi) + \psi$ if $\neg\psi \notin K * \phi$	$\phi \supset \psi \leq \neg\phi$ if $\phi \supset (\psi \wedge \chi) \leq \neg\phi$ i.e., if $\psi \wedge \chi \notin K * \phi$	$\gamma(\llbracket\phi\rrbracket) \cap \llbracket\psi\rrbracket$ if $\gamma(\llbracket\phi\rrbracket) \cap \llbracket\psi\rrbracket \neq \emptyset$	$\{S \cup (\llbracket K\rrbracket \cap \llbracket\phi\rrbracket) : S \in \$_F$ or $S = \emptyset\}$ if $\llbracket K\rrbracket \cap \llbracket\phi\rrbracket \neq \emptyset$
	$F * \psi$ if $\neg\psi \in K * \phi$	$\psi \leq_F \chi$ otherwise	$\gamma_F(\llbracket\psi\rrbracket)$ otherwise	$\{S \cup (S^F_\phi \cap \llbracket\phi\rrbracket) : S \in \$_F$ or $S = \emptyset\}$ otherwise
Radical	$K * (\phi \wedge \psi)$ [if $\neg\psi \notin Cn(\phi)$ $Cn(\psi)$ otherwise]	$\phi \supset \psi \leq \phi \supset \chi$ [and: if $\vdash \psi$ then $\vdash \chi$]	$\gamma(\llbracket\phi \wedge \psi\rrbracket)$ [if $\llbracket\phi \wedge \psi\rrbracket \neq \emptyset$ $\llbracket\psi\rrbracket$ otherwise]	$\{S \cap \llbracket\phi\rrbracket : S \in \$$ and $S \cap \llbracket\phi\rrbracket \neq \emptyset\}$ $[\cup \{W\}]$
Conservative	$(K * \phi) + \psi$ if $\neg\psi \notin K * \phi$	$\phi \supset \psi \leq \neg\phi$ if $\phi \supset (\psi \wedge \chi) \leq \neg\phi$ i.e., if $\psi \wedge \chi \notin K * \phi$	$\gamma(\llbracket\phi\rrbracket) \cap \llbracket\psi\rrbracket$ if $\gamma(\llbracket\phi\rrbracket) \cap \llbracket\psi\rrbracket \neq \emptyset$	$\{S \cup (S_\phi \cap \llbracket\phi\rrbracket) :$ $S \in \$$ or $S = \emptyset\}$
	$K * \psi$ if $\neg\psi \in K * \phi$	$\psi \leq \chi$ otherwise	$\gamma(\llbracket\psi\rrbracket)$ otherwise	
Moderate	$K * (\phi \wedge \psi)$ if $\neg\psi \notin Cn(\phi)$	$\phi \supset \psi \leq \phi \supset \chi$ if $\psi \wedge \chi \notin Cn(\phi)$	$\gamma(\llbracket\phi \wedge \psi\rrbracket)$ if $\llbracket\phi \wedge \psi\rrbracket \neq \emptyset$	$\{S \cap \llbracket\phi\rrbracket : S \in \$$ and $S \cap \llbracket\phi\rrbracket \neq \emptyset\}$
	$K * \psi$ if $\neg\psi \in Cn(\phi)$	$\psi \leq \chi$ otherwise	$\gamma(\llbracket\psi\rrbracket)$ otherwise	$\cup\ \{S \cup \llbracket\phi\rrbracket : S \in \$$ or $S = \emptyset\}$

Table 1: Four methods for iterated belief change

slogan '*Preserve as much of your belief state as possible.*' Belief states, identified with unary iterated revision functions in part I of this paper, can be identified with belief-revision guiding structures, *provided that* a method for changing belief-revision guiding structures is agreed upon. We focussed on entrenchment relations as belief-revision guiding structures, but briefly hinted at the possibility of alternatively taking semantic choice functions or systems of spheres.

We implemented the conservative idea in the rule for changing entrenchment relations. Not surprisingly, this rule entails AGM's c-conservatism. But this strengthened form of conservatism has undesirable effects as well. These effects were already noted by Darwiche and Pearl in their reaction to Boutilier's similar idea of implementing conservativity in belief revision in the mid-1990s.

The construction in this paper has shown, however, that surprisingly, the project of a more far-reaching conservatism than AGM can be pursued in a context without any dispositional coherence, i.e., without the axioms $(*7)$ and $(*8)$ which I take to be the landmark axioms of AGM.

At this place it is appropriate to add an explanation of why getting rid of $(*7)$ and $(*8)$ is indeed a relevant. The dispositional condition $(*8)$ is notoriously violated by all revision operations for belief sets that are generated through changes of belief bases, in the following sense. If the belief set K (which is logically closed) is generated from the belief base H (which is not logically closed and for which $Cn(H) = K$), then ψ is defined to be in $K * \phi$ if it follows from all maximal subsets of H consistent with ϕ. Consider the base $H = \{\neg p, \neg q, p \supset q\}$. Then it is easy to check that – as in the model-theoretic example just discussed – $K * (p \vee q) = Cn(\{p \leftrightarrow \neg q\})$ and $K * p = Cn(\{p\})$, which again violates $(*8)$.

To mention another important example, if we have a revision function generated by minimization with respect to a partial (not total) ordering $\preceq$ of interpretations, then $(*8)$ is no longer valid. The idea in this model is to set $K = th(min_{\preceq} W)$ and $K * \phi = th(min_{\preceq}\{w \in W : w \text{ satisfies } \phi\})$, for every ϕ.[32]

Now consider, for example, the pair $\langle W, \preceq \rangle$ consisting of the set W of interpretations $w_0 = \langle \neg p, \neg q \rangle$, $w_1 = \langle p, \neg q \rangle$, $w_2 = \langle \neg p, q \rangle$ and $w_3 = \langle p, q \rangle$, together with an ordering $\preceq$ on W such that $w_0 \preceq w_i$ for $i = 1, 2, 3$ and $w_2 \preceq w_3$, but that these are the only pairs related by the ordering $\preceq$. Clearly $\preceq$ is not total. Then we get $K*(p\vee q) = th(\{w_1, w_2\}) = Cn(\{p \leftrightarrow \neg q\})$ and $K*p = th(\{w_1, w_3\}) = Cn(\{p\})$, so clearly $K*((p\vee q)\wedge p) = K*p \subseteq (K*(p\vee q))+p$. But not $(K*(p\vee q))+p \subseteq K*p$, although $\neg p \notin K * (p \vee q)$. Thus $(*8)$ is violated.[33]

[32] Here W is the set of all possible worlds, $th(W')$ for some set W' of possible worlds it the set true in every world in W'. This modelling is used, for instance, in Rott (1993). Boutilier's (1993, 1996) model is essentially the same, but he does require completeness of the ordering relation.

[33] It is no mere coincidence that this examples is similar to the first one using belief bases. See Lewis (1981), Nebel (1989) and Rott (1993).

Another good example against ($*8$) is intuitively discussed in the contexts of counterfactuals by Ginsberg (1986) and adapted for a belief revision context in Rott (2001, pp. 188–191).

Striking counterexamples against the dispositional condition ($*7$) are much harder to come by. I have recently given one in Rott (200).

In this paper we have shown that the specific problems of iterated belief change don't require ($*7$) and ($*8$), nor are any weaker substitutes for them necessary. Even on the basis of synchronic coherence, diachronic coherence is independent of dispositional coherence. So much for the good news. On the negative side, however, we have to put up with the fact that even the criticism of Darwiche and Pearl can be reproduced in this much more general context. Conservative belief revision exhibits a capricious evaluation of the recency of information and it thus violates what can be regarded as a fourth kind of cohrence: temporal coherence.

Surveying the alternatives that have been offered, we discussed what I called external, radical and moderate belief revision. The last method has been discussed in quite a number of papers, it has sneaked into the folklore of the field without there being a standard reference paper. Nevertheless a consensus seems to be emerging, confirmed by the findings of this paper, that this method give the best results one can expect from a non-numerical approach. A doxastic agent should not be too conservative, and she should not be too radical either; as Aristotle knew, "a master of any art avoids excess and defect, but seeks the intermediate and chooses this – the intermediate".[34]

The approach follows the strategy of giving priority to new information over old information thoughout a whole sequence of inputs. It does so at the expense of conservatism. Not only does it contradict the strong conservatism with respect to belief-revision guiding structures. We have shown that not even the modest conservativity of AGM in revisions that are not belief-contravening is valid any more. We conclude that very little is left of the idea of conservatism in belief change theories, and we do not see any reason why this should be regretted.

It is remarkable, though, that like conservative revision, moderate revision is feasible without the aid of dispositional coherence.

Acknowledgement

I would like to thank Carlos Areces, Johan van Benthem, Michael Freund, David Makinson, Abhaya Nayak, Wolfgang Spohn, and other people from audiences in Amsterdam, Konstanz and Paris for helpful comments on early versions of this paper.

[34] *Nicomachean Ethics*, Book II, 1106b, transl. W.D. Ross.

Appendix: Proofs

Proof of Observation 7.

With the exception of (X-Minimality) and (X-Representation) where there are direct proofs, we show that $\leq_X$ has a certain property on the assumption that $\leq$ has it.

(Reflexivity) $\phi \leq_X \phi$ holds, since clearly either $\phi \notin X$ or both $\phi \in X$ and $\phi \leq \phi$, by the reflexivity of $\leq$.

(Extensionality) Let $\phi \dashv\vdash \psi$. Then, $\phi \in X$ iff $\psi \in X$, since X is a belief set. So $\phi \leq_X \chi$ if and only if $\psi \leq_X \chi$, since we know that $\phi \leq \chi$ if and only if $\psi \leq \chi$, by the extensionality of $\leq$. Similarly, $\chi \leq_X \phi$ if and only if $\chi \leq_X \psi$, since we know that $\chi \leq \phi$ if and only if $\chi \leq \psi$, again by the extensionality of $\leq$.

(Choice) In order to show that $\phi \wedge \psi \leq_X \chi$ if and only if $\phi \leq_X \psi \wedge \chi$ or $\psi \leq_X \phi \wedge \chi$, we have to show that

$$\phi \wedge \psi \notin X, \text{ or } \chi \in X \text{ and } \phi \wedge \psi \leq \chi \quad (\dagger)$$

is equivalent with

$$\phi \notin X, \text{ or } \psi \wedge \chi \in X \text{ and } \phi \leq \psi \wedge \chi, \text{ or } \psi \notin X, \text{ or } \phi \wedge \chi \in X \text{ and } \psi \leq \phi \wedge \chi (\ddagger)$$

($\dagger$) implies ($\ddagger$): First suppose that $\phi \wedge \psi \notin X$. Then either $\phi \notin X$ or $\psi \notin X$, since X is a belief set, and we are done. So suppose that $\phi \wedge \psi \in X$, and that $\chi \in X$ and $\phi \wedge \psi \leq \chi$. Then $\phi \wedge \chi \in X$ and $\psi \wedge \chi \in X$, since X is a belief set. From $\phi \wedge \psi \leq \chi$, we get that either $\phi \leq \psi \wedge \chi$ or $\psi \leq \phi \wedge \chi$, since we assume that $\leq$ satisfies Choice, so in either case one disjunct of ($\ddagger$) is satisfied.

($\ddagger$) implies ($\dagger$): If $\phi \notin X$ or $\psi \notin X$, then $\phi \wedge \psi \notin X$, since X is a belief set, and we are done. So suppose that $\psi \wedge \chi \in X$ and $\phi \leq \psi \wedge \chi$. Then $\chi \in X$, since X is a belief set, and $\phi \wedge \psi \leq \chi$ by Choice for $\leq$, and we are done. The last case, $\phi \wedge \chi \in X$ and $\psi \leq \phi \wedge \chi$ is nalogous.

(Maximality) Suppose that $\top \leq_X \phi$. Then either $\top \notin X$ or both $\phi \in X$ and $\top \leq \phi$. But $\top \notin X$ is impossible since X is a belief set. So $\top \leq \phi$, from which we conclude that $\vdash \phi$, by Maximality for $\leq$.

(X-Minimality) By the definition of $\leq_X$, it follows immediately from $\phi \notin X$ that $\phi \leq_X \psi$ for all ψ.

(X-Representation) Suppose that $\phi \in X$ and that $\phi \leq_X \psi$. Then it follows immediately from the definition of $\leq_X$ that $\psi \in X$.

(Transitivity) Suppose that $\phi \leq_X \psi$ and $\psi \leq_X \chi$. The former means that $\phi \notin X$ or both $\psi \in X$ and $\phi \leq \psi$. If $\phi \notin X$, then we immediately get $\phi \leq_X \chi$, as desired. If, however, $\psi \in X$ and $\phi \leq \psi$, then we conclude from $\psi \in X$ and $\psi \leq_X \chi$ that

$\chi \in X$ and $\psi \leq \chi$. Since $\leq$ is assumed to be transitive, we get that $\phi \leq \chi$ and thus $\phi \leq_X \chi$, as desired. $\Box$

Proof of Observation 8 (Soundness for conservative belief change).

Let $\leq$ determine the iterated revision function $*$ by the conservative method (CER).

First observe that (CER) conforms to the Triangle Property: Since $\phi \supset (\psi \wedge \bot)$, by the Extensionality and Reflexivity of $\leq$, (CER) immediately gives us $\psi \leq^{*C}_{\phi} \bot$ if and only if $\phi \supset \psi \leq \neg\phi$. Since the triangle property holds, we can freely use both (Def $K_{\leq}$) and (Def$^{\omega}*$ from $\leq$) for the construction of revised belief sets.

Since for all Φ, $\leq^{*C}_{\Phi}$ is a basic entrenchment relation and $K * (\Phi \cdot \psi)$ is determined by $\leq^{*C}_{\Phi}$ via (Def$^{\omega}*$ from $\leq$), we know that $(*1^{\omega})$, $(*2^{\omega})$, (5^{ω}) and $(*6^{\omega})$ are satisfied, by Observation 6, part (a).

By parts (b) and (c) of the same Observation and by Observation 7, we also know that $\leq_{\Phi}$ is faithful to $K * \Phi$, and that $(*3^{\omega})$ and $(*4^{\omega})$ are satisfied for all $K * \Phi$, with the possible exception of the case $n = 0$. But the case $n = 0$ is trivial since $\leq$ satisfies K-Minimality and K-Representation by the definition of $*$'s being determined by $\leq$.

We now prove that belief change determined by moderate entrenchment revision (CER) satisfies $(*9\mathrm{C}^{\omega})$:

$$\text{If } \neg\chi \in K * (\Phi \cdot \psi), \text{ then } K * (\Phi \cdot \psi \cdot \chi) = K * (\Phi \cdot \chi)$$

First, the limiting case $\bot \in Cn(\chi)$ yields $\mathcal{L}$ as resulting belief set on the left-hand side and on the right-hand side, by the definition of an entrenchment relation determining an iterated revision function. So let for the rest of this proof $\bot \notin Cn(\chi)$.

Second, let $\neg\chi \in K * (\Phi \cdot \psi)$. By (Def $K_{\leq}$), this means that not $\neg\chi \leq^{*C}_{\Phi\cdot\psi} \bot$.

Limiting case. $\neg\psi \in Cn(\emptyset)$. Then the definition of conservative entrenchment revision prescribes that $\leq^{*C}_{\Phi\cdot\psi} = \leq^{*C}_{\Phi}$, so $\leq^{*C}_{\Phi\cdot\psi\cdot\chi} = \leq^{*C}_{\Phi\cdot\chi}$ and $K * (\Phi \cdot \psi \cdot \chi) = K * (\Phi \cdot \chi)$.

Principal case. $\neg\psi \notin Cn(\emptyset)$. Then, by (CER), not $\neg\chi \leq^{*C}_{\Phi\cdot\psi} \bot$ means that not $\psi \supset \neg\chi \leq^{*C}_{\Phi} \neg\psi$, or equivalently, $\neg\psi <^{*C}_{\Phi} \psi \supset \neg\chi$ $(*)$.

Under the premises of the prinicipal case, we reason as follows:

$\xi \in K * (\Phi \cdot \psi \cdot \chi)$ iff (by Def $K_{\leq}$)

not $\xi \leq^{*C}_{\Phi\cdot\psi\cdot\chi} \bot$ iff (by CER and $\chi \not\vdash \bot$)

$$\text{not } \left\{ \begin{array}{ll} \chi \supset \xi \leq^{*C}_{\Phi\cdot\psi} \neg\chi & \text{if } \chi \supset (\bot \wedge \xi) \leq^{*C}_{\Phi\cdot\psi} \neg\chi \\ \xi \leq^{*C}_{\Phi\cdot\psi} \bot & \text{if } \neg\chi <^{*C}_{\Phi\cdot\psi} \chi \supset (\bot \wedge \xi) \end{array} \right\} \text{ iff (since } \leq^{*C}_{\Phi\cdot\psi} \text{ is reflexive)}$$

not $\chi \supset \xi \leq^{*C}_{\Phi\cdot\psi} \neg\chi$ iff (by CER)

$$\text{not} \left\{ \begin{array}{ll} \psi \supset (\chi \supset \xi) \leq^{*C}_{\Phi} \neg\psi & \text{if } \psi \supset ((\chi \supset \xi) \wedge \neg\chi) \leq^{*C}_{\Phi} \neg\psi \\ \chi \supset \xi \leq^{*C}_{\Phi} \neg\chi & \text{if } \neg\psi <^{*C}_{\Phi} \psi \supset ((\chi \supset \xi) \wedge \neg\chi) \end{array} \right\}$$

iff (by extensionality of $\leq^{*C}_{\Phi}$ and $(*)$)

not $\chi \supset \xi \leq^{*C}_{\Phi} \neg\chi$ iff (since $\leq^{*C}_{\Phi}$ is reflexive)

$$\text{not} \left\{ \begin{array}{ll} \chi \supset \xi \leq^{*C}_{\Phi} \neg\chi & \text{if } \chi \supset (\bot \wedge \xi) \leq^{*C}_{\Phi} \neg\chi \\ \xi \leq^{*C}_{\Phi} \bot & \text{if } \neg\chi <^{*C}_{\Phi} \chi \supset (\bot \wedge \xi) \end{array} \right\} \quad \text{iff (by CER)}$$

not $\xi \leq^{*C}_{\Phi\cdot\chi} \bot$ (by Def $K_{\leq}$)

$\xi \in K * (\Phi \cdot \chi)$.

Since we started with the left-hand side of the identity in $(*9\text{C}^{\omega})$ and we finished with its right-hand side, we are done. □

Proof of Observation 10 (Completeness for conservative belief change).

Let $*$ satisfy $(*1^{\omega})$ – $(*6^{\omega})$ as well as $(*9\text{C}^{\omega})$. We write K for $*(\mathbf{0})$. We define the initial entrenchment relation $\leq \; = \; \leq^{*C}_{\mathbf{0}}$ by using the idea of (Def $\leq$ from $*$) and putting

$$\phi \leq \psi \quad \text{iff} \quad \phi \notin *(\langle \neg(\phi \wedge \psi) \rangle) \text{ or } \vdash \psi$$

By Observation 2, $\leq$ is a basic entrenchment relation, and by Observation 1 (applied to $\phi = \top$) we get that $K * \top = \{\phi : \bot < \phi\} = K_{\leq}$. Since $*$ satisfies $(*3^{\omega})$ and $(*4^{\omega})$, $K = K * \top$, if K is consistent. So in this case $K = K_{\leq}$. But K might also be inconsistent. Either way, $\leq$ is faithful to K, i.e. $\leq$ satisfies K-Minimality and K-Representation, and K is as required for the definition of conservative determination.

By Lemma 3 and Observation 7, all entrenchment relations that are generated from a basic entrenchment relation by conservative entrenchment revision satisfy Connectedness.

Having taken care of the case $\Phi = \mathbf{0}$, we now need to show that for all $\Phi = \langle \phi_1, \ldots, \phi_n$, $n \geq 1$, it holds that

$$*(\Phi) = \left\{ \begin{array}{ll} K_{\leq^{*C}_{\Phi}} = \{\psi : \bot <^{*C}_{\Phi} \psi\} & \text{if } \nvdash \neg\phi_n \\ \mathcal{L} & \text{if } \vdash \neg\phi_n \end{array} \right.$$

The lower line follows immediately from $(*5^{\omega})$. The upper line will be shown by induction.

Induction basis $n = 1$. We have to show that

$$*(\phi) = \left\{ \begin{array}{ll} K_{\leq^{*C}_{\phi}} = \{\psi : \bot <^{*C}_{\phi} \psi\} & \text{if } \nvdash \neg\phi \\ \mathcal{L} & \text{if } \vdash \neg\phi \end{array} \right.$$

By the Triangle Property, the crucial condition $\bot <^{*C}_{\phi} \psi$ is equivalent to $\neg\phi < \phi \supset \psi$. So the above condition reduces precisely to the idea of (Def $*$ from $\leq$). Since $\leq$ was retrieved from $*$ with the help of (Def $\leq$ from $*$), the claim for the induction basis is exactly the content of Observation 1.

Inductive step. Suppose the claim holds for all Φ of length n ($n \geq 1$; let Φ–1 be an abbreviation for $\langle\phi_1 \cdot \ldots \cdot \phi_{n-1}\rangle$). Does the claim hold for $\Phi \cdot \phi_{n+1}$, too?

We have to show that

$$*(\Phi \cdot \phi_{n+1}) = \begin{cases} \{\psi : \bot <^{*C}_{\Phi\cdot\phi_{n+1}} \psi\} & \text{if } \nvdash \neg\phi_{n+1} \\ \mathcal{L} & \text{if } \vdash \neg\phi_{n+1} \end{cases}$$

The lower line is trivial. So suppose that $\nvdash \neg\phi_{n+1}$.

$\psi \in *(\Phi \cdot \phi_{n+1})$ iff (by ($*$3,4,9C$^\omega$))

$$\begin{cases} \neg\phi_{n+1} \notin K * \Phi & \text{and } \phi_{n+1} \supset \psi \in K * \Phi \quad \textit{or} \\ \neg\phi_{n+1} \in K * \Phi & \text{and } \psi \in K * ((\Phi - 1) \cdot \phi_{n+1}) \end{cases}$$

iff (by inductive hypothesis, $\nvdash \neg\phi_n$, $\nvdash \neg\phi_{n-1}$)

$$\begin{cases} \neg\phi_{n+1} \leq^{*C}_{\Phi} \bot & \text{and } \bot <^{*C}_{\Phi} \phi_{n+1} \supset \psi \quad \textit{or} \\ \bot <^{*C}_{\Phi} \neg\phi_{n+1} & \text{and } \bot <^{*C}_{(\Phi-1)\cdot\phi_{n+1}} \psi \end{cases}$$

iff (by the triangle property)

$$\begin{cases} \phi_n \supset \neg\phi_{n+1} \leq^{*C}_{\Phi-1} \neg\phi_n & \text{and } \neg\phi_n <^{*C}_{\Phi-1} \phi_n \supset (\phi_{n+1} \supset \psi) \quad \textit{or} \\ \neg\phi_n <^{*C}_{\Phi-1} \phi_n \supset \neg\phi_{n+1} & \text{and } \neg\phi_{n+1} <^{*C}_{\Phi-1} \phi_{n+1} \supset \psi \end{cases}$$

iff (by logic)

$$\text{not: } \begin{cases} \phi_n \supset \neg\phi_{n+1} \leq^{*C}_{\Phi-1} \neg\phi_n & \text{and } \phi_n \supset (\phi_{n+1} \supset \psi) \leq^{*C}_{\Phi-1} \neg\phi_n \quad \textit{or} \\ \neg\phi_n <^{*C}_{\Phi-1} \phi_n \supset \neg\phi_{n+1} & \text{and } \phi_{n+1} \supset \psi \leq^{*C}_{\Phi-1} \neg\phi_{n+1} \end{cases}$$

iff (by CER)

not: $\phi_{n+1} \supset \psi \leq^{*C}_{\Phi} \neg\phi_{n+1}$

iff (by connectedness)

$\neg\phi_{n+1} <^{*C}_{\Phi} \phi_{n+1} \supset \psi$

iff (by the triangle property, $\nvdash \neg\phi_{n+1}$)

$\bot <^{*C}_{\Phi\cdot\phi_{n+1}} \psi$. $\Box$

Proof of Observation 12 (Darwiche-Pearl postulates).

Since for each Φ, $\leq^{*C}_{\Phi}$ is an entrenchment relation that satisfies Minimality and Representation with respect to $K * \Phi$, we may without loss of generality suppose that $\Phi = \mathbf{0}$ in each of the Darwiche-Pearl postulates.

Now let $*$ satisfy $(*1^{\omega})$ – $(*6^{\omega})$ and $(*9\text{C}^{\omega})$. According to $(*3,4,9\text{C})$, we have

$$(\dagger) \qquad \xi \in (K * \psi) * \chi \;\text{ iff }\; \begin{cases} \xi \in K * \chi \text{ and } \neg\chi \in K * \psi & \textit{or} \\ \chi \supset \xi \in K * \psi \text{ and } \neg\chi \notin K * \psi \end{cases}$$

(a) For $(\text{DP}2^{\omega})$, let $\chi \vdash \neg\psi$. This is equivalent to $\psi \vdash \neg\chi$. We have to show that ξ is in $K * \chi$ if and only if it is in $(K * \psi) * \chi$, that is, by $(\dagger)$,

$$(\ddagger) \qquad \xi \in K * \chi \;\text{ iff }\; \begin{cases} (\text{I}) \quad \xi \in K * \chi \text{ and } \neg\chi \in K * \psi & \textit{or} \\ (\text{II}) \quad \chi \supset \xi \in K * \psi \text{ and } \neg\chi \notin K * \psi \end{cases}$$

From left to right. Assume that $\xi \in K * \chi$. Then (I) can fail only if $\neg\chi \notin K * \psi$. In this case we are left to show for (II) that $\chi \supset \xi \in K * \psi$. Since $\psi \vdash \neg\chi$, we get from $(*2)$ and $(*1)$ that $\neg\chi \in K * \psi$, so by $(*1)$ again, we get $\chi \supset \xi \in K * \psi$.

From right to left. First note that (II) is impossible since we just verified that $\psi \vdash \neg\chi$ entails $\chi \supset \xi \in K * \psi$. But if (I) holds, then the LHS follows immediately.

For $(\text{DP}3^{\omega})$, let $\psi \in K * \chi$. We have to show that $\psi \in (K * \psi) * \chi$, that is, by $(\dagger)$ that either (I) $\psi \in K * \chi$ and $\neg\chi \in K * \psi$, or (II) $\chi \supset \psi \in K * \psi$ and $\neg\chi \notin K * \psi$. But by the supposition, (I) can only fail if $\neg\chi \notin K * \psi$. In this case we are left to show for (II) that $\chi \supset \psi \in K * \psi$. This follows immediately from $(*2)$ and $(*1)$.

For $(\text{DP}4^{\omega})$, let $\neg\psi \notin K * \chi$. We have to show that $\neg\psi \notin (K * \psi) * \chi$, that is, by $(\dagger)$ that both (I) $\neg\psi \notin K * \chi$ or $\neg\chi \notin K * \psi$, and (II) $\chi \supset \neg\psi \notin K * \psi$ or $\neg\chi \in K * \psi$. But (I) follows immediately from the supposition. And (II) is true, since by $(*2)$ $\psi \in K * \psi$, so if it is the case that $\chi \supset \neg\psi \in K * \psi$, then we get with the help of $(*1)$ that $\neg\chi \in K * \psi$.

(b) Now suppose that $*$ in addition satisfies $(*7^{\omega})$ and $(*8^{\omega})$. For $(\text{DP}1^{\omega})$, let $\chi \vdash \psi$. We have to show that $(\ddagger)$ holds.

From left to right. Assume that $\xi \in K * \chi$. Then (I) can fail only if $\neg\chi \notin K * \psi$. In this case we are left to show for (II) that $\chi \supset \xi \in K * \psi$. But by $\chi \vdash \psi$ and $(*6)$, $K * \chi = K * (\psi \wedge \chi)$, so $\xi \in K * (\psi \wedge \chi)$. By $(*7)$, then $\xi \in Cn((K * \psi) \cup \{\chi\})$. By $(*1)$, this means that $\chi \supset \xi \in K * \psi$, as desired.

From right to left. If (I) holds, then the LHS follows immediately. So suppose that (II) holds, i.e., that $\chi \supset \xi \in K * \psi$ and $\neg\chi \notin K * \psi$. From the latter and $(*8)$, we get that $Cn((K * \psi) \cup \{\chi\}) \subseteq K * (\psi \wedge \chi)$. From the former, we get that $\xi \in Cn((K * \psi) \cup \{\chi\})$. Hence, $\xi \in K * (\psi \wedge \chi)$, and by $\chi \vdash \psi$ and $(*6)$, this means that $\xi \in K * \chi$, as desired. $\Box$

Proof of Observation 13.

(Reflexivity) and (Extensionality) of $\leq_\phi^{*M}$ follow directly from the Reflexivity and Extensionality of $\leq$ and (MER).

(Choice) We have to show that $\psi \wedge \chi \leq_\phi^{*M} \xi$ iff $\psi \leq_\phi^{*M} \chi \wedge \xi$ or $\chi \leq_\phi^{*M} \psi \wedge \xi$. By (MER), this follows directly from (Choice) for $\leq$ if $\phi \vdash \psi \wedge \chi \wedge \xi$. If, however, $\phi \nvdash \psi \wedge \chi \wedge \xi$, the claim reduces to $\phi \supset (\psi \wedge \chi) \leq \phi \supset \xi$ iff $\phi \supset \psi \leq \phi \supset (\chi \wedge \xi)$ or $\phi \supset \chi \leq \phi \supset (\psi \wedge \xi)$. But this follows from Extensionality and Choice for $\leq$.

(Maximality). We have to show that $\top \leq_\phi^{*M} \psi$ implies that $\psi \in Cn(\emptyset)$. By (MER), this follows directly from Maximality for $\leq$ if $\phi \vdash \top \wedge \psi$. So suppose that $\phi \nvdash \top \wedge \psi$, i.e., $\psi \nvdash \psi$. By (MER), then $\top \leq_\phi^{*M} \psi$ means that $\phi \supset \top \leq \phi \supset \psi$. By (Maximality) for $\leq$, then $\phi \supset \psi \in Cn(\emptyset)$. But this contradicts $\psi \nvdash \psi$.

Now let $\leq$ be a basic entrenchment relation.
(Transitivity). We need to show that

$$\psi \leq_\phi^{*M} \chi \text{ and } \chi \leq_\phi^{*M} \xi \text{ together imply } \psi \leq_\phi^{*M} \xi$$

The antecedents of this implication amount to

(i) If $\phi \vdash \psi \wedge \chi$ then $\psi \leq \chi$
(ii) If $\phi \nvdash \psi \wedge \chi$ then $\phi \supset \psi \leq \phi \supset \chi$
(iii) If $\phi \vdash \chi \wedge \xi$ then $\chi \leq \xi$
(iv) If $\phi \nvdash \chi \wedge \xi$ then $\phi \supset \chi \leq \phi \supset \xi$

Now we are considering cases.

Case 1. Let $\phi \vdash \psi \wedge \xi$. We need to show that $\psi \leq \xi$.

Case 1a. Suppose that also $\phi \vdash \chi$. Then the claim follows from (i) and (iii) and the transitivity of $\leq$.

Case 1b. Suppose that also $\phi \nvdash \chi$. Then $\phi \nvdash \psi \wedge \chi$, so by (ii) $\phi \supset \psi \leq \phi \supset \chi$. Since $\vdash \phi \supset \psi$, Maximality for $\leq$ also gives us $\vdash \phi \supset \chi$, contradicting $\phi \nvdash \chi$.

Case 2. Let $\phi \nvdash \psi \wedge \xi$. We need to show that $\phi \supset \psi \leq \phi \supset \xi$.

Case 2a. Suppose that both $\phi \nvdash \psi$ and $\phi \nvdash \xi$. Then $\phi \nvdash \psi \wedge \chi$ and $\phi \nvdash \chi \wedge \xi$. Then the claim follows from (ii) and (iv) and the transitivity of $\leq$.

Case 2b. Suppose that also $\phi \vdash \xi$. Then $\vdash \phi \supset \xi$. So by Extensionality, Choice and Reflexivity for $\leq$, we get $\phi \supset \psi \leq \phi \supset \xi$.

Case 2c. Suppose that also $\phi \vdash \psi$. Then $\phi \nvdash \xi$. Hence $\phi \nvdash \chi \wedge \xi$. By, (iv), we get that $\phi \supset \chi \leq \phi \supset \xi$. Now first suppose that $\vdash \phi \supset \chi$, then by Maximality for $\leq$ also $\vdash \phi \supset \xi$, and we get a contradiction with $\phi \nvdash \xi$. So suppose secondly that

$\not\vdash \phi \supset \chi$. So by (ii) $\phi \supset \psi \leq \phi \supset \chi$, and the claim follows from the transitivity of $\leq$. $\Box$

Proof of Observation 14 (Soundness for moderate belief change).

Let $\leq$ determine the iterated revision function $*$ by the moderate method (MER).

First observe that (MER) conforms to the Triangle Property: Let $\neg\phi \notin Cn(\emptyset)$. Since $\psi \wedge \bot \notin Cn(\phi)$, (MER) gives us $\psi \leq^{*M}_{\phi} \bot$ if and only if $\phi \supset \psi \leq \phi \supset \bot$, which is equivalent to $\phi \supset \psi \leq \neg\phi$, by the Extensionality of $\leq$. Since the triangle property holds, we can freely use both (Def $K_{\leq}$) and (Def$^{\omega}*$ from $\leq$) for the construction of revised belief sets.

Since for all Φ, $\leq^{*M}_{\Phi}$ is a basic entrenchment relation by Observation 13, and $K * (\Phi \cdot \psi)$ is determined by $\leq^{*M}_{\Phi}$ via (Def$^{\omega}*$ from $\leq$), we know that ($*1^{\omega}$), ($*2^{\omega}$), (5^{ω}) and ($*6^{\omega}$) are satisfied, by Observation 6, part (a).

We now prove that belief change determined by moderate entrenchment revision (MER) satisfies ($*9\mathrm{M}^{\omega}$):

$$K * (\Phi \cdot \psi \cdot \chi) = \begin{cases} K * (\Phi \cdot (\psi \wedge \chi)) & \text{if } \neg\chi \in Cn(\psi) \\ K * (\Phi \cdot \chi) & \text{otherwise} \end{cases}$$

First, the limiting case $\bot \in Cn(\chi)$ yields $\mathcal{L}$ as resulting belief set on the left-hand side and on the right-hand side, by the definition of an entrenchment relation determining an iterated revision function. So let for the rest of this proof $\bot \notin Cn(\chi)$.

Now we reason as follows:

$\xi \in K * (\Phi \cdot \psi \cdot \chi)$ iff (by Def $K_{\leq}$)

$\bot <^{*M}_{\Phi \cdot \psi \cdot \chi} \xi$ iff (by MER and $\chi \not\vdash \bot$)

$\chi \supset \bot <^{*M}_{\Phi \cdot \psi} \chi \supset \xi$ iff (by MER)

$\left\{ \begin{array}{ll} \psi \supset (\chi \supset \bot) <^{*M}_{\Phi} \psi \supset (\chi \supset \xi) & \text{if } \psi \not\vdash (\chi \supset \bot) \wedge (\chi \supset \xi) \\ \chi \supset \bot <^{*M}_{\Phi} \chi \supset \xi & \text{if } \psi \vdash (\chi \supset \bot) \wedge (\chi \supset \xi) \end{array} \right\}$ iff (by prop.logic)

$\left\{ \begin{array}{ll} (\psi \wedge \chi) \supset \bot <^{*M}_{\Phi} (\psi \wedge \chi) \supset \xi & \text{if } \psi \wedge \chi \not\vdash \bot \\ \chi \supset \bot <^{*M}_{\Phi} \chi \supset \xi & \text{if } \psi \vdash \neg\chi \end{array} \right\}$ iff (by MER and $\chi \not\vdash \bot$)

$\left\{ \begin{array}{ll} \bot <^{*M}_{\Phi \cdot (\psi \wedge \chi)} \xi & \text{if } \psi \not\vdash \neg\chi \\ \bot <^{*M}_{\Phi \cdot \chi} \xi & \text{if } \psi \vdash \neg\chi \end{array} \right\}$ iff (by Def $K_{\leq}$)

$\left\{ \begin{array}{ll} \xi \in K * (\Phi \cdot (\psi \wedge \chi)) & \text{if } \psi \not\vdash \neg\chi \\ \xi \in K * (\Phi \cdot \chi)\xi & \text{if } \psi \vdash \neg\chi \end{array} \right\}$.

Since we started with the left-hand side and finished with the right-hand side of $(*9\mathrm{M}^\omega)$, we are done.

Now we give a counterexample to all the other AGM conditions.

Let $K = Cn(\{\neg p, \neg q\})$, and let the one-step revision function $*$ defined as follows. $K * \phi = K + \phi$ for every ϕ consistent with K, and $K * \phi = Cn(\phi)$ for ϕ identical with one of $\pm p \wedge \pm q$, where $\pm$ is either nothing or $\neg$. For the rest, let us have $K * p = Cn(p \wedge \neg q)$, $K * q = Cn(q)$, $K * (p \leftrightarrow \neg q) = Cn(p \leftrightarrow \neg q)$ and $K * (p \vee q) = Cn(p \wedge q)$.

Obviously, this is allowed by $(*1)$, $(*2)$, $(*5)$ and $(*6)$, which are the only conditions required for one-step revisions determined by basic entrenchment, according to Observations 1 and 2.

For the iterated case, we may now use $(*9\mathrm{M}^\omega)$, the validity of which we have just shown. Since $\neg p \notin Cn(p \vee q)$, we get $K * (p \vee q) * p = K * ((p \vee q) \wedge p) = K * p$ $= Cn(p \wedge \neg q)$.

Now it is easy to see that we have a violation of AGM's postulates $(*3^\omega)$ and $(*4^\omega)$ with respect to $K * (p \vee q)$:

Re $(*3^\omega)$: We find that $Cn(p \wedge \neg q) = K * (p \vee q) * p \ \not\subseteq\ (K * (p \vee q)) + p = (Cn(p \wedge q)) + p = Cn(p \wedge q)$.

Re $(*4^\omega)$: And we find that $\neg p \notin K * (p \vee q) = Cn(p \wedge q)$, but still $Cn(p \wedge q) = K * (p \vee q) \ \not\subseteq\ K * (p \vee q) * p = Cn(p \wedge \neg q)$.

The very same example can serve as a counterexample to $(*7^\omega)$ and $(*8^\omega)$ with respect to K. □

Proof of Observation 16 (Completeness for moderate belief change).

Let $*$ satisfy $(*1^\omega)$ – $(*2^\omega)$ and $(*5^\omega)$ – $(*6^\omega)$ as well as $(*9\mathrm{M}^\omega)$. We write K for $*(\mathbf{0})$. We define the initial entrenchment relation $\leq\, = \,\leq^{*C}_{\mathbf{0}}$ by using the idea of (Def $\leq$ from $*$) and putting

$$\phi \leq \psi \quad \text{iff} \quad \phi \notin *(\langle \neg(\phi \wedge \psi) \rangle) \text{ or } \vdash \psi$$

By Observation 13 and Lemma 3, all entrenchment relations that are generated from a basic entrenchment relation by moderate entrenchment revision satisfy Connectedness.

Without $(*3^\omega)$ and $(*4^\omega)$, K need not be related to any revised belief set, and $\leq$ need not be faithful to K. But we need to show that for all $\Phi = \langle \phi_1, \ldots, \phi_n \rangle$, $n \geq 1$, it holds that

$$*(\Phi) = \begin{cases} K_{\leq^{*M}_{\Phi}} = \{\psi : \bot <^{*M}_{\Phi} \psi\} & \text{if } \nvdash \neg\phi_n \\ \mathcal{L} & \text{if } \vdash \neg\phi_n \end{cases}$$

The lower line follows immediately from $(*5^\omega)$. The upper line will be shown by induction.

Induction basis $n = 1$. We have to show that

$$*(\phi) = \begin{cases} K_{\leq^{*M}_\phi} = \{\psi : \bot <^{*M}_\phi \psi\} & \text{if } \not\vdash \neg\phi \\ \mathcal{L} & \text{if } \vdash \neg\phi \end{cases}$$

By the Triangle Property, the crucial condition $\bot <^{*M}_\phi \psi$ is equivalent to $\neg\phi < \phi \supset \psi$. So the above condition reduces precisely to the idea of (Def $*$ from $\leq$). Since $\leq$ was retrieved from $*$ with the help of (Def $\leq$ from $*$), the claim for the induction basis is exactly the content of Observation 1.

Inductive step. Suppose the claim holds for all Φ of length n ($n \geq 1$; let Φ–1 be an abbreviation for $\langle\phi_1 \cdot \ldots \cdot \phi_{n-1}\rangle$). Does the claim hold for $\Phi \cdot \phi_{n+1}$, too?

We have to show that

$$*(\Phi \cdot \phi_{n+1}) = \begin{cases} \{\psi : \bot <^{*M}_{\Phi\cdot\phi_{n+1}} \psi\} & \text{if } \not\vdash \neg\phi_{n+1} \\ \mathcal{L} & \text{if } \vdash \neg\phi_{n+1} \end{cases}$$

The lower line is trivial. So suppose that $\not\vdash \neg\phi_{n+1}$.

$\psi \in *(\Phi \cdot \phi_{n+1})$ iff (by $(*3{,}4{,}9\mathrm{M}^\omega)$)

$$\begin{cases} \neg\phi_{n+1} \notin Cn(\phi_n) & \text{and } \psi \in K * ((\Phi - 1) \cdot (\phi_n \wedge \phi_{n+1})) \quad \textit{or} \\ \neg\phi_{n+1} \in Cn(\phi_n) & \text{and } \psi \in K * ((\Phi - 1) \cdot \phi_{n+1}) \end{cases}$$

iff (by inductive hypothesis, $\not\vdash \neg\phi_n$, $\not\vdash \neg\phi_{n-1}$)

$$\begin{cases} \neg\phi_{n+1} \notin Cn(\phi_n) & \text{and } \bot <^{*M}_{(\Phi-1)\cdot(\phi_n\wedge\phi_{n+1})} \psi \quad \textit{or} \\ \neg\phi_{n+1} \in Cn(\phi_n) & \text{and } \bot <^{*M}_{(\Phi-1)\cdot\phi_{n+1}} \psi \end{cases}$$

iff (by the triangle property)

$$\begin{cases} \neg\phi_{n+1} \notin Cn(\phi_n) & \text{and } \neg(\phi_n \wedge \phi_{n+1}) <^{*M}_{\Phi-1} (\phi_n \wedge \phi_{n+1}) \supset \psi \quad \textit{or} \\ \neg\phi_{n+1} \in Cn(\phi_n) & \text{and } \neg\phi_{n+1} <^{*M}_{\Phi-1} \phi_{n+1} \supset \psi \end{cases}$$

iff (by logic)

$$\text{not: } \begin{cases} \neg\phi_{n+1} \notin Cn(\phi_n) & \text{and } (\phi_n \wedge \phi_{n+1}) \supset \psi \leq^{*M}_{\Phi-1} \neg(\phi_n \wedge \phi_{n+1}) \quad \textit{or} \\ \neg\phi_{n+1} \in Cn(\phi_n) & \text{and } \phi_{n+1} \supset \psi \leq^{*M}_{\Phi-1} \neg\phi_{n+1} \end{cases}$$

iff (by logic)

$$\text{not: } \begin{cases} (\phi_{n+1} \supset \psi) \wedge \neg\phi_{n+1} \notin Cn(\phi_n) & \text{and } \phi_n \supset (\phi_{n+1} \supset \psi) \leq^{*M}_{\Phi-1} \phi_n \supset \neg\phi_{n+1} \quad \textit{or} \\ (\phi_{n+1} \supset \psi) \wedge \neg\phi_{n+1} \in Cn(\phi_n) & \text{and } \phi_{n+1} \supset \psi \leq^{*M}_{\Phi-1} \neg\phi_{n+1} \end{cases}$$

iff (by MER)

not: $\phi_{n+1} \supset \psi \leq^{*M}_{\Phi} \neg\phi_{n+1}$
iff (by connectedness)

$\neg\phi_{n+1} <^{*M}_{\Phi} \phi_{n+1} \supset \psi$
iff (by the triangle property, $\not\vdash \neg\phi_{n+1}$)

$\bot <^{*M}_{\Phi\cdot\phi_{n+1}} \psi$. $\Box$

References

Alchourrón, Carlos, Peter Gärdenfors and David Makinson: 1985, 'On the Logic of Theory Change: Partial Meet Contraction Functions and Their Associated Revision Functions', *Journal of Symbolic Logic* **50**, 510–530.

Alchourrón, Carlos, and David Makinson: 1985, 'On the Logic of Theory Change: Safe Contraction', *Studia Logica* **44**, 405–422.

Areces, Carlos, and Veronica Becher: 2001, 'Iterable AGM Functions', in Mary-Anne Williams and Hans Rott (eds.), *Frontiers of Belief Revision*, Dordrecht: Kluwer, 261–277.

Boutilier, Craig: 1993, 'Revision Sequences and Nested Conditionals', in R. Bajcsy (ed.), *IJCAI-93 – Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 519–525.

Boutilier, Craig: 1996, 'Iterated Revision and Minimal Change of Conditional Beliefs', *Journal of Philosohical Logic* **25**, 263–305.

Cantwell, John: 1999, 'Some Logics of Iterated Belief Change', *Studia Logica* **63**, 49–84.

Darwiche, Adnan, and Judea Pearl: 1994, 'On the Logic of Iterated Belief Revision', in Ronald Fagin, ed., *TARK'94 – Proceedings of the Fifth Conference on Theoretical Aspects of Reasoning About Knowledge*, Pacific Grove, Cal.: Morgan Kaufmann, pp. 5–23.

Darwiche, Adnan, and Judea Pearl: 1997, 'On the Logic of Iterated Belief Revision', *Artificial Intelligence* **89**, 1–29.

Fermé, Eduardo: 2000, 'Irrevocable Belief Revision and Epistemic Entrenchment', *Logic Journal of the IGPL* **8**, 645–652.

Freund, Michael, and Daniel Lehmann: 1994, 'Belief Revision and Rational Inference', Technical Report TR-94-16, Institute of Computer Science, Hebrew University, Jerusalem.

Friedman, Nir, and Joseyph Y. Halpern: 1996, 'Belief Revision: A Critique', in L. C. Aiello, J. Doyle and S. C. Shapiro (eds.), *Principles of Knowledge Representation and Reasoning. Proceedings of the Fifth International Conference (KR'96)*, Morgan Kaufmann, San Mateo, Cal., 421–431.

Friedman, Nir, and Joseyph Y. Halpern: 1999, 'Belief Revision: A Critique', *Journal of Logic, Language and Information* **8**, 401–420.

Gärdenfors, Peter: 1988, *Knowledge in Flux. Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press, Cambridge, Mass.

Gärdenfors, Peter, and David Makinson: 1988, 'Revisions of Knowledge Systems Using Epistemic Entrenchment', in Moshe Vardi (ed.), *TARK'88 – Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, Los Altos: Morgan Kaufmann, pp. 83–95.

Gärdenfors, Peter, and Hans Rott: 1995, 'Belief revision', in D. M. Gabbay, C. J. Hogger, and J. A. Robinson (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming Volume IV: Epistemic and Temporal Reasoning*, Oxford University Press, pp. 35–132.

Ginsberg, Matthew L.: 1986, 'Counterfactuals', *Artificial Intelligence* **30**, 35–79.

Glaister, Stephen: 1998, 'Symmetry and Belief Revision', *Erkenntnis* **49**, pp. 21-56.

Grove, Adam: 1988, 'Two Modellings for Theory Change', *Journal of Philosophical Logic* **17**, 157–170.

Hansson, Sven O.: 1992, 'A Dyadic Representation of Belief', in P. Gärdenfors (ed.), *Belief Revision*, Cambridge: Cambridge University Press, pp. 89–121.

Hansson, Sven O.: 1993, 'Reversing the Levi Identity', *Journal of Philosophical Logic*, **22**, 637–669.

Hansson, Sven O.: 1994, 'Kernel Contraction', *Journal of Symbolic Logic*, **59**, 845–859.

Hansson, Sven O.: 1999, *A Textbook of Belief Dynamics: Theory Change and Database Updating*, Dordrecht: Kluwer.

Kelly, Kevin: 1999, 'Iterated Belief Revision, Reliability, and Inductive Amnesia', *Erkenntnis* **50**, 11–58.

Konieczny, Sébastien: 1998, 'Operators With Memory for Iterated Revision', Université de Lille I, Technical Report No LIFL IT-314, May 1998.

Konieczny, Sébastien, and Ramón Pino Pérez: 2000, 'A Framework for Iterated Revision', *Journal of Applied Non-Classical Logic* **10**, 339–367.

Lehmann, Daniel: 1995, 'Belief Revision, Revised', in *IJCAI'95 – Proceedings of the 14th International Joint Conference on Artificial Intelligence*, San Mateo: Morgan Kaufmann, pp. 1534–1540.

Lewis, David: 1981, "Ordering semantics and premise semantics for counterfactuals", *Journal of Philosophical Logic* **10**, 217–234.

Liberatore, Paolo: 1997, 'The Complexity of Iterated Belief Revision', in *Proceedings of the Sixth International Conference on Database Theory (ICDT'97)*, Lecture Notes in Computer Science **1186**, Berlin: Springer, pp. 276–290.

Lindström, Sten, and Wlodzimierz Rabinowicz: 1992, 'Belief Revision, Epistemic Conditionals and the Ramsey Test', *Synthese* **91**, 195–237.

Nayak, Abhaya C.: 1994, 'Iterated Belief Change Based on Epistemic Entrenchment', *Erkenntnis* **41**, 353–390.

Nayak, Abhaya C., Norman Y. Foo, Maurice Pagnucco and Abdul Sattar: 1996, 'Changing Conditional Beliefs Unconditionally', in Yoav Shoham (ed.), *TARK'96 – Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge*, Morgan Kaufmann, pp. 119–135.

Nayak, Abhaya C., Paul Nelson and Hanan Polansky: 1996, 'Belief Change as Change in Epistemic Entrenchment', *Synthese* **109**, 143–174.

Nebel, Bernhard: 1989, "A knowledge level analysis of belief revision", in R. Brachman, H. Levesque and R. Reiter (eds.), *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, Ca., pp. 301–311.

Pagnucco, Maurice, and Hans Rott: 1999, 'Severe Withdrawal (and Recovery)', *Journal of Philosophical Logic* **28**, 501–547.

Papini, Odile: 2001, 'Iterated Revision Operations Stemming from the History of an Agent's Observations', in Mary-Anne Williams and Hans Rott (eds.), *Frontiers of Belief Revision*, Dordrecht: Kluwer, 279–301.

Rott, Hans: 1991, 'Two Methods of Constructing Contractions and Revisions of Knowledge Systems', *Journal of Philosophical Logic* **20**, 149–73.

Rott, Hans: 1991, 'A Non-monotonic Conditional Logic for Belief Revision I', in André Fuhrmann and Michael Morreau (eds.), *The Logic of Theory Change*, Lecture Notes in Computer Science **465**, Berlin: Springer, pp. 135–181.

Rott, Hans: 1992, 'Preferential Belief Change Using Generalized Epistemic Entrenchment', *Journal of Logic, Language and Information* **1**, 45–78.

Rott, Hans: 1993, 'Belief Contraction in the Context of the General Theory of Rational Choice', *Journal of Symbolic Logic* **58**, pp. 1426–1450.

Rott, Hans: 1998, 'Logic and Choice,' in Itzhak Gilboa (ed.), *TARK'98 – Proceedings of the Seventh Conference on Theoretical Aspects of Rationality and Knowledge*, San Francisco: Morgan Kaufmann, pp. 235–248.

Rott, Hans: 1999, 'Coherence and Conservatism in the Dynamics of Belief. Part I: Finding the Right Framework', *Erkenntnis* **50**, 387–412.

Rott, Hans: 2000, 'Two Dogmas of Belief Revision', *Journal of Philosophy* **97**, 503–522.

Rott, Hans: 2001, *Change, Choice and Inference*, Oxford University Press.

Rott, Hans: 2001, 'Basic Entrenchment', to appear in *Studia Logica*.

Rott, Hans: 2002, 'A Counterexample to Three Fundamental Principles of Belief Formation', manuscript Regensburg, March 2002.

Schlechta, Karl: 1991, 'Some Results on Theory Revision', in André Fuhrmann and Michael Morreau (eds.), *The Logic of Theory Change*, Lecture Notes in Computer Science **465**, Springer, Berlin: Springer, pp. 72-92.

Segerberg, Krister: 1998, 'Irrevocable Belief Revision in Dynamic Doxastic Logic', *Notre Dame Journal of Formal Logic*, **39**, pp. 287–306.

Spohn, Wolfgang: 1988, 'Ordinal Conditional Functions', in William L. Harper and Brian Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics*, Vol. II, Dordrecht: Reidel, pp. 105–134.

Williams, Mary-Anne: 1994, 'Transmutations of Knowledge Systems', in Jon Doyle, Erik Sandewall, and Pierro Torasso (eds.), *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann Publishers, pp. 619–629.

Department of Philosophy
University of Regensburg
93040 Regensburg
Germany

Appendix II (not to be published): An example for moderate revision without dispositional coherence

Consider the language with two propositional variables p and q.

Let $w_0 = \langle \overline{p}, \overline{q} \rangle$, $w_1 = \langle p, \overline{q} \rangle$, $w_2 = \langle \overline{p}, q \rangle$ and $w_3 = \langle p, q \rangle$.

The semantics of entrenchment according to Rott (2001, 251) says that

$\phi < \psi$ iff $]\!]\phi[\![\cap \gamma(\,]\!]\{\phi, \psi\}[\![\,) \neq \emptyset$ and $]\!]\psi[\![\cap \gamma(\,]\!]\{\phi, \psi\}[\![\,) = \emptyset$

Now we take a choice function that violates (I), (III) and (II) (i.e., Sen's properties α, β and γ).

$$\begin{aligned}
\gamma(\{w_0, \dots\}) &= \{w_0\} \\
\gamma(\{w_i\}) &= \{w_i\} \text{ for } i = 1,2,3,4 \\
\gamma(\{w_1, w_2, w_3\}) &= \{w_3\} \\
\gamma(\{w_1, w_2\}) &= \{w_1, w_2\} \\
\gamma(\{w_1, w_3\}) &= \{w_1\} \\
\gamma(\{w_2, w_3\}) &= \{w_2, w_3\}
\end{aligned}$$

The problematic part of the example is the pair $\gamma(\{w_1, w_2, w_3\})$ and $\gamma(\{w_1, w_3\})$. It is not completely unrealistic; the form corresponds to my "Counterexample … ". It violates choice conditions (I) and (III), and together with $\gamma(\{w_1, w_2\})$ it violates (II). Let us first rewrite the last four lines of the above:

$$\begin{aligned}
\gamma(\,]\!]\neg p \wedge \neg q[\![\,) &= \{\langle p, q \rangle\} \\
\gamma(\,]\!]p \leftrightarrow q[\![\,) &= \,]\!]p \leftrightarrow q[\![\\
\gamma(\,]\!]\neg p[\![\,) &= \{\langle p, \overline{q} \rangle\} \\
\gamma(\,]\!]\neg q[\![\,) &= \,]\!]\neg q[\![
\end{aligned}$$

From the first of these four lines, we get $\neg p \leq \neg q$, $\neg q \leq \neg p$,
$\neg p < p \leftrightarrow q$, $\neg p \vee \neg q < p \leftrightarrow q$,
$\neg q < p \leftrightarrow q$,
$\neg q < p \rightarrow q$,
$\neg p < q \rightarrow p$.

From the second line, we get $p \rightarrow q \leq q \rightarrow p$ and $q \rightarrow p \leq p \rightarrow q$.

From the third line, we get $p \rightarrow q < \neg p \vee \neg q$.

From the fourth line, we get $q \rightarrow p \leq \neg p \vee \neg q$ and $\neg p \vee \neg q \leq q \rightarrow p$.

Using the above definition of entrenchments from semantic choice functions, we get the following table:

	$\neg p \wedge \neg q$	$\neg p$	$p \leftrightarrow q$	$\neg q$	$p \rightarrow q$	$\neg p \vee \neg q$	$q \rightarrow p$	$\top$
$\neg p \wedge \neg q$		$\leq$	$<$	$\leq$	$<$	$\leq$	$<$	$<$
$\neg p$	$\leq$		$<$	$\leq$	$\leq$	$<$	$<$	$<$
$p \leftrightarrow q$	$>$	$>$		$>$	$\leq$	$>$	$\leq$	$<$
$\neg q$	$\leq$	$\leq$	$<$		$<$	$\leq$	$\leq$	$<$
$p \rightarrow q$	$>$	$\leq$	$\leq$	$>$		$<$	$\leq$	$<$
$\neg p \vee \neg q$	$\leq$	$>$	$<$	$\leq$	$>$		$\leq$	$<$
$q \rightarrow p$	$>$	$>$	$\leq$	$\leq$	$\leq$	$\leq$		$<$
$\top$	$>$	$>$	$>$	$>$	$>$	$>$	$>$	

(Notice that, e.g., the set $\{\phi : \neg p < \phi\}$ or the set $\{\phi : \neg q < \phi\}$ is not a theory!)

Semantically, using γ, one can see that the following are the results of potential revisions:

$$\begin{aligned} K &= Cn(\neg p \wedge \neg q) \\ K * p &= Cn(p \wedge \neg q) \\ K * q &= Cn(q) \\ K * (p \leftrightarrow \neg q) &= Cn(p \leftrightarrow \neg q) \\ K * (p \vee q) &= Cn(p \wedge q) \end{aligned}$$

Using $(*9\mathrm{M})$, we have

$K * (p \vee q) * p =$(mod.rev., $\neg p \notin Cn(p \vee q)$)

$K * ((p \vee q) \wedge p) =$

$K * p = Cn(p \wedge \neg q)$.

In this example, we have a violation of AGM's third and fourth postulate with respect to $K * (p \vee q)$:

Re $(*3^{\omega})$
$Cn(p \wedge \neg q) = K * (p \vee q) * p \not\subseteq (K * (p \vee q)) + p = (Cn(p \wedge q)) + p = Cn(p \wedge q)$.

Re $(*4^{\omega})$
We have $\neg p \notin K * (p \vee q) = Cn(p \wedge q)$, but still
$Cn(p \wedge q) = K * (p \vee q) \not\subseteq K * (p \vee q) * p = Cn(p \wedge \neg q)$.

Now let us construct the revision of $\leq$ by $p \vee q$.
Given $\psi \wedge \chi \notin Cn(p \vee q)$, we get $\psi \leq^*_{p\vee q} \chi$ iff $(p \vee q) \supset \psi \leq (p \vee q) \supset \chi$.
In particular, $\bot <^*_{p\vee q} \chi$ iff $\neg(p \vee q) < \neg(p \vee q) \vee \chi$.

Looking in the table above, we find that the following sentences in K satisfy this last condition (and thus make up, as it were, $K \dot{-} \neg(p \vee q)$):
$p \leftrightarrow q$, $p \rightarrow q$, $q \rightarrow p$, $\top$, that is, exactly $Cn(p \leftrightarrow q)$.

$K_{\leq^*_{p\vee q}} = \{\psi : \bot <^*_{p\vee q} \psi\} =$
$\{\psi : (p \vee q) \supset \psi \in Cn(p \leftrightarrow q)\} =$
$(Cn(p \leftrightarrow q)) + (p \vee q) = Cn(p \wedge q)$.

By the first line of (MER), we get for all ψ in K except $\top$,
$\psi \leq^*_{p\vee q} (p \vee q)$,
so $p \vee q$ is ranked highly in the revised belief set (as desired).

For $\top$, the second line of (MER) applies and gives us $(p \vee q) \leq^*_{p\vee q} \top$ (as desired).

Now let us perform a second revision, this time by p. Using nothing but (MER), we want to find the condition for $\psi \ (\leq^*_{p\vee q})^*_p \ \chi$

Case 1. $p \nvdash \psi \wedge \chi$. Then we have $p \supset \psi \ \leq^*_{p\vee q} \ p \supset \chi$

Case 1a. $p \vee q \nvdash (p \supset \psi) \wedge (p \supset \chi)$. This means that $p \nvdash \psi \wedge \chi$ (so the condition is the same as the one for Case 1). And we get a reduction to $(p \vee q) \supset (p \supset \psi) \leq (p \vee q) \supset (p \supset \chi)$ or equivalently $p \supset \psi \leq p \supset \chi$

Case 1b. $p \vee q \vdash (p \supset \psi) \wedge (p \supset \chi)$. This is inconsistent with Case 1.

Case 2. $p \vdash \psi \wedge \chi$. Then we have $\psi \ \leq^*_{p\vee q} \ \chi$

Case 2a. $p \vee q \nvdash \psi \wedge \chi$. Then we get a reduction to $(p \vee q) \supset \psi \leq (p \vee q) \supset \chi$

Case 2b. $p \vee q \vdash \psi \wedge \chi$. Then we get a reduction to $\psi \leq \chi$

In sum, then, we have

$$\psi \ (\leq^*_{p\vee q})^*_p \ \chi \ \text{ iff } \ \begin{cases} p \supset \psi \leq p \supset \chi & \text{if } \ p \nvdash \psi \wedge \chi \\ (p \vee q) \supset \psi \leq (p \vee q) \supset \chi & \text{if } \ p \vdash \psi \wedge \chi, \text{ but } p \vee q \nvdash \psi \wedge \chi \\ \psi \leq \chi & \text{if } \ p \vee q \vdash \psi \wedge \chi \end{cases}$$

Returning now to our example, let us check for the belief set $K_{\leq^*_{(p\vee q)\cdot p}}$.

First, is $p \wedge q$ in this belief set, i.e., do we not get $p \wedge q \leq^*_{(p\vee q)\cdot p} \bot$?
This problem falls under Case 1 of the above, so the question reduces to $p \supset (p \wedge q) \leq p \supset \bot$, or equivalently, to $p \supset q \leq \neg p$. This is in fact true (see the big table), so $p \wedge q \notin K_{\leq^*_{(p\vee q)\cdot p}}$, as predicted by the use of (∗9M).

Second, is $p \wedge \neg q$ in this belief set, i.e., do we not get $p \wedge \neg q \leq^*_{(p \vee q) \cdot p} \bot$?
This problem, too, falls under Case 1 of the above, so the question reduces to $p \supset (p \wedge \neg q) \leq p \supset \bot$, or equivalently, to $\neg p \vee \neg q \leq \neg p$. This is not true (see the big table), so $p \wedge \neg q \in K_{\leq^*_{(p \vee q) \cdot p}}$, as predicted by the use of (*9M).

This result formally confirms in terms of entrenchment revision (MER) what we said about the failure of (*3) and (*4) by virtue of (*9).

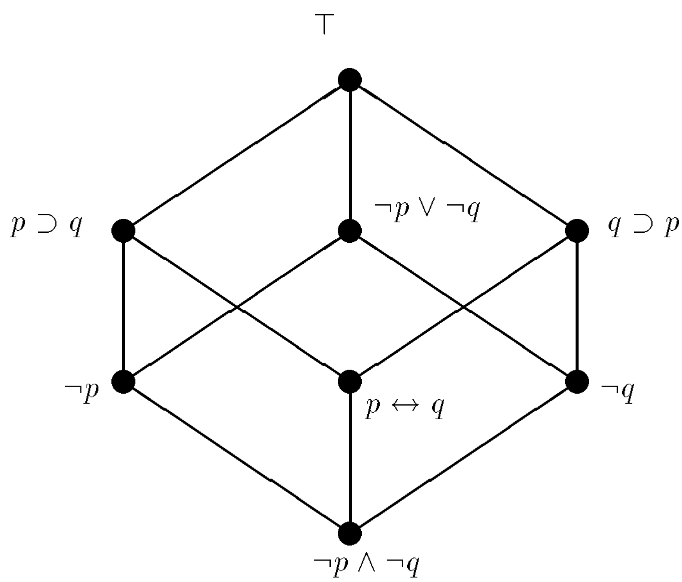

Figure 2: *Hasse diagram for* $K = Cn(\{\neg p, \neg q\})$