Fuhrmann, A. and Morreau M. (eds.) *The Logic of Theory Change*,
Springer Lecture Notes in Artificial Intelligence 465, 1991

EPISTEMIC ENTRENCHMENT WITH INCOMPARA-BILITIES AND RELATIONAL BELIEF REVISION*

Sten Lindström and Wlodzimierz Rabinowicz[1]

1. Introduction

In an earlier paper (Lindström & Rabinowicz, 1989), we proposed a generalization of the approach to belief revision due to Alchourrón-Gärdenfors-Makinson. Our proposal was to view belief revision as a relation rather than as a function on theories (or belief sets). The idea was to allow for there being several equally reasonable revisions of a theory with a given proposition. The same approach was used — and developed further — in our study (1990) of belief revision and epistemic conditionals.

In the present paper, we want to show that the relational approach is the natural result of generalizing in a certain way an approach to belief revision which is due to Adam Grove. In his (1988) paper, Grove presents two closely related modelings of functional belief revision, one in terms of a family of "spheres" around the agent's theory G and the other in terms of an epistemic entrenchment ordering of propositions.[2] Intuitively, a proposition A is at least as entrenched in the agent's belief set as another proposition B if and only if the following holds: provided the agent would have to revise his beliefs so as to falsify the conjunction $A \wedge B$, he should do it in such a way as to allow for the falsity of B.

The "sphere"-terminology is natural when one looks upon theories and propositions as being represented by sets of possible worlds. Grove's spheres may be thought of as possible "fallback" theories relative to the agent's original theory: theories that he may reach by deleting propositions that are not "sufficiently" entrenched (according to standards of sufficient entrenchment of varying stringency). To put it differently, fallbacks are theories that are closed upwards under entrenchment: if T is a fallback, A belongs to T and B is at least as entrenched as A, then B also belongs to T. The entrenchment ordering can be recovered from the family of fallbacks by the definition: A is at least as entrenched as B iff A belongs to every fallback to which B belongs.

Representing theories and propositions as sets of possible worlds, the following picture illustrates Grove's family of spheres around a given theory G and his definition of revision. Notice that the spheres around a theory are "nested", i.e., simply ordered. For any two spheres, one is included in the other. Grove's family of spheres closely resembles Lewis' sphere semantics for counterfactuals, the main difference being that Lewis' spheres are "centered" around a single world instead of a theory (a set of worlds).

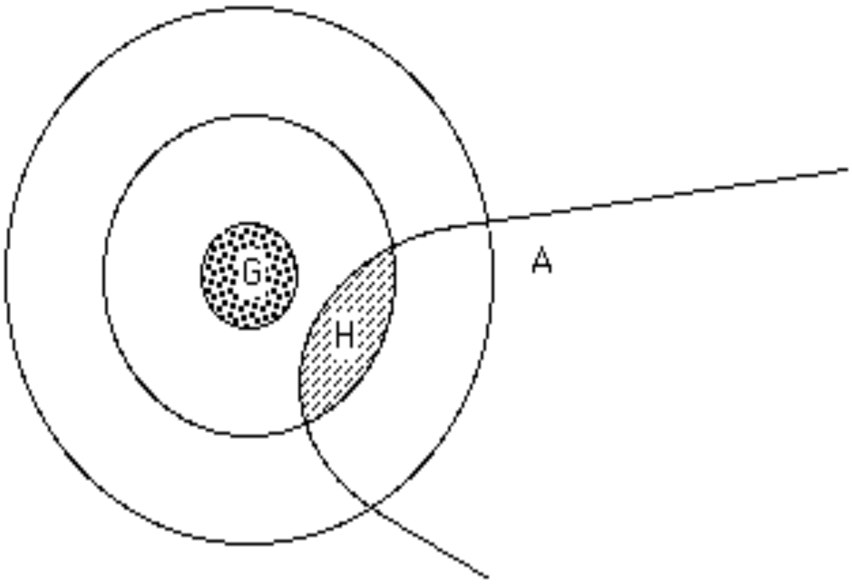

Figure 1.

The shaded area H in Figure 1 represents the revision of G with a proposition A. The revision of G with A is defined as the strongest A-permitting fallback theory of G expanded with A. In the possible worlds representation, this is the intersection of A with the smallest sphere around G that is compatible with A. (Any revision has to contain the proposition we revise with. Therefore, if A is logically inconsistent, the revision with A is taken to be the inconsistent theory.)

The relational notion of belief revision that we are interested in, results from weakening epistemic entrenchment by not assuming it to be *connected.* In other words, we want to allow that some propositions may be incomparable with respect to epistemic entrenchment. As a result, the family of fallbacks around a given theory will no longer have to be nested. It will no longer be a family of spheres but rather a family of "ellipses". This change opens up the possibility for several different ways of revising a theory with a given proposition.

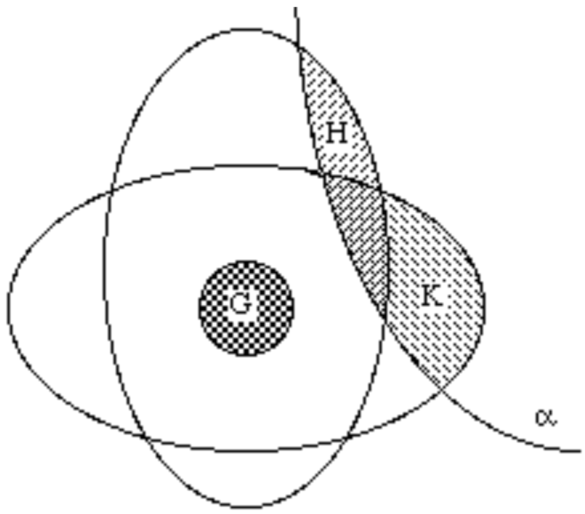

Figure 2.

In this figure, the two ellipses represent two different fallback theories for G, each of which is a strongest A-permitting fallback. Consequently, there are two possible revisions of G with A: each one of H and K is the intersection of A with a strongest A-permitting fallback.

Given connectedness for epistemic entrenchment, there is a way of defining revision directly from entrenchment, without going via fallbacks. The revision of G with A (in symbols, G^*_A) may then be defined thus (for the interesting case in which A is logically consistent):

$B \in G^*_A$ iff $A \rightarrow B$ is more entrenched than $A \rightarrow \neg B$ (relative to G).

G^*_A defined in this way will coincide with the intersection of A with the maximal A-permitting fallback. However, when we give up connectedness there may not be a *unique* fallback of this kind. Therefore, this direct way of defining revision from entrenchment is no longer available.

Why is it that a non-nested family of fallbacks allows for incomparabilities with respect to epistemic entrenchment? Figure 3 below provides an explanation. In this picture, there is one fallback theory, H, in which A is true but B is not, and another, K, in which B is true but A is not. Hence, neither B is at least as entrenched as A nor vice versa. A and B are incomparable.

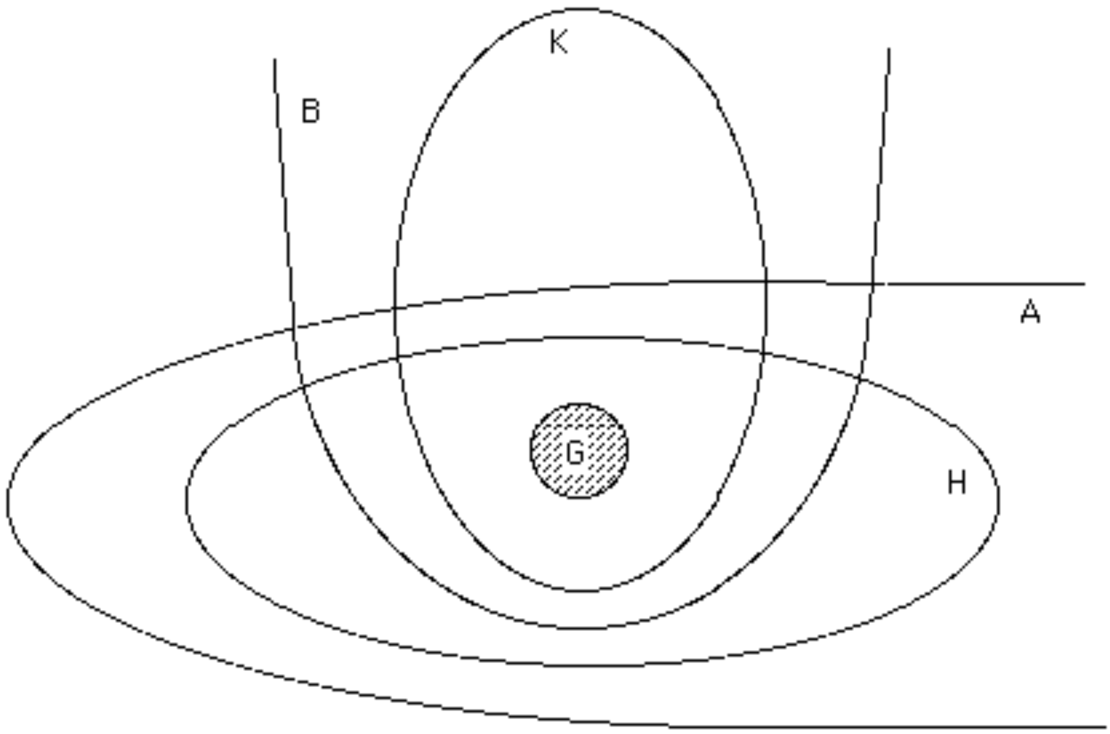

Figure 3.

The next section provides the logical background for our discussion. In Section 3, we introduce an axiom system for epistemic entrenchment without connectedness. We also provide an axiomatic characterization of fallback families. We prove that the two notions are interdefinable and that, in fact, there is a one-to-one correspondence between entrenchment orderings and fallback families.

In Section 3, we also show that non-connected epistemic entrenchment orderings can be viewed as intersections of families of connected entrenchment orderings. There is an analogy here with probabilistic representations of epistemic states. Working with a non-connected entrenchment ordering is analogous to a representation of the agent's epistemic state by an "interval-valued" probability assignment. On the other hand, the representation in terms of a family of connected entrenchment orderings is like the approach in which an epistemic state is viewed as a class of real-valued probability assignments. Gärdenfors' functional approach in terms of a unique connected entrenchment ordering corresponds to the classical Bayesian approach in which the epistemic state is identified with a unique real-valued probability assignment. Of course, this is only an analogy; entrenchment orderings and probability functions do not obey the same axioms.

In Section 4, we define relational belief revision from (non-connected) epistemic entrenchment, via the intermediate notion of a fallback family. We also show that the original entrenchment ordering can be recovered from the resulting belief revision relation by means of the following definition:

($D_{\leq}$) A is at least as entrenched in G as B iff either (i) A is logically true or (ii) for every theory H, if H is a possible revision of G with $\neg(A \wedge B)$, then B does not belong to H.[3]

In the same section, we also provide an axiom system for relational belief revision, which, given the definition ($D_{\leq}$), is sufficient for the derivation of our postulates for epistemic entrenchment. We also prove that our axiom system for belief revision is not sufficiently strong to guarantee that all belief revision systems that satisfy our axioms are *representable*, that is, that they are definable in the appropriate way from an epistemic entrenchment relation. In the special case, when we consider functional belief revision only, our axioms are sufficient to guarantee representability. This result is essentially equivalent to Grove's representation theorem for Gärdenfors-type belief revision in terms of families of spheres. However, the problem of giving a nice axiomatic characterization of just those relational belief revision systems that are representable still remains to be solved.

Finally, in Section 5, we discuss the relationship between belief revision and belief *contraction.* According to Levi (1980), belief revision can be seen as a two-step process: a contraction followed by an expansion. In order to revise our theory G with A, we first contract G with ¬A — move to a weaker theory K in which ¬A is no longer present — and then we expand K with A (i.e., we add A to K and close under logical consequence). In our approach, revision is also viewed as a two-step process, but the first step is different. Instead of contracting G with A, we move to a maximal A-permitting fallback of G. In Section 5, we argue that contractions and maximal fallbacks are different concepts. Nor is it possible to construct contraction from fallbacks, or what amounts to the same, from epistemic entrenchment. Such a construction would yield the so-called postulate of Recovery for contraction, which we reject. According to this postulate, if $A \in G$ and K is a contraction of G with A, one can recover the original theory G from K by expanding the latter with A.

On the other hand, it is possible to impose reasonable constraints on contraction in the presence of which Levi's definition of revision in terms of contraction will follow from our definition in terms of fallbacks. In this way, the two approaches may, after all, be complementary.

2. Logics

We consider a fixed sentential language $\mathcal{L}$ with the following symbols: (i) atomic sentences: p_1, p_2,... ; (ii) classical connectives: $\bot$ (falsity), $\rightarrow$ (the material conditional); (iii) possibly some non-classical connectives; (iv) parentheses. The set Φ of *sentences* is the smallest set such that: (i) all atomic sentences are in Φ; (ii) $\bot \in \Phi$; (iii) if A, B $\in \Phi$, then $(A \rightarrow B)$ is in Φ; (iv) if $\Box$ is an n-ary non-classical connective in $\mathcal{L}$ and A_1,..., A_n are in Φ, then $\Box(A_1,..., A_n) \in \Phi$. The connectives $\neg$ (negation), $\wedge$ (conjunction), $\vee$ (disjunction) and $\leftrightarrow$ (the material biconditional) are introduced as abbreviations in the metalanguage in the usual way. For instance, $\neg A = A \rightarrow \bot$.

A *logic* L is a set of sentences such that: (i) all truth-functional tautologies are in L; (ii) L is closed under *modus ponens*. For every logic L and set Γ of sentences, we define $Cn_L(\Gamma)$ to be the smallest set of sentences that includes $\Gamma \cup L$ and is closed under modus ponens. We say that A is an *L-consequence* of Γ (also written as $\Gamma \vdash_L A$) if $A \in Cn_L(\Gamma)$. Note that every logic L satisfies *compactness* and the *deduction theorem:*

(i) if $\Gamma \vdash_L A$, then for some finite $\Delta \subseteq \Gamma$, $\Delta \vdash_L A$;

(ii) if $\Gamma \cup \{A\} \vdash_L B$, then $\Gamma \vdash_L A \rightarrow B$.

Let L be a logic. L is (absolutely) *consistent* if and only if $L \neq \Phi$. A set $\Gamma \subseteq \Phi$ is said to be *L-consistent* if $Cn_L(\Gamma) \neq \Phi$. A is said to be L-consistent if {A} is L-consistent. Γ is an *L-theory* if and only if $L \subseteq \Gamma$ and Γ is closed under modus ponens. In other words, Γ is an L-theory just in case $\Gamma = Cn_L(\Gamma)$. Note that every L-theory is a logic in its own right. A set Γ is *L-maximal* if and only if Γ is L-consistent and for every Δ, if $\Gamma \subseteq \Delta$ and Δ is L-consistent, then $\Gamma = \Delta$. Every L-maximal set is an L-theory.

G, H, K, T, T′,... are variables ranging over L-theories. The set of such theories is denoted by $\mathcal{T}_L$.

3. Epistemic Entrenchment

In the following, we let L be a consistent logic.

Definition 3.1. Let G be any L-theory. An *epistemic entrenchment ordering for* G is a binary relation $\leq$ of the sentences in $\mathcal{L}$ satisfying the following postulates:

($\leq$1) If $A \leq B$ and $B \leq C$, then $A \leq C$. *(Transitivity)*

($\leq$2) If $A \vdash_L B$, then $A \leq B$. *(Dominance)*

($\leq$3) If $A \leq B$ and $A \leq C$, then $A \leq B \wedge C$. *(Conjunctive Closure)*

($\leq$4) If $\bot \notin G$, then $A \notin G$ iff $A \leq \bot$. *(Bottom)*

($\leq$5) If $\top \leq A$, then $\vdash_L A$. *(Top)*

Here, $T =_{df} \neg\bot$.

If $A \leq B$, we say that B *is at least as (epistemically) entrenched in* G *as* A. (If we want to make the reference to G explicit, we write $A \leq_G B$ rather than $A \leq B$. Sometimes, we might also write $B \geq A$ or $B \geq_G A$ instead of $A \leq B$.) The corresponding strict ordering < is defined by: $A < B$ iff $A \leq B$ and not $B \leq A$. $A < B$ is to be read as "B *is (epistemologically) more entrenched in* G *than* A". We also define the corresponding equality relation, A *and* B *are (epistemologically) equally entrenched in* G (in symbols, $A \approx B$ or $A \approx_G B$), as: $A \leq B$ and $B \leq A$.

Let us now look at the intuitive motivation for the above axioms. The theory G consists of all those sentences (in $\mathcal{L}$) that are accepted (or believed) by the agent. Although the agent has full belief in all the sentences in G, some of his beliefs are held more firmly than others: the more firm — or epistemically entrenched — a belief is, the less willing the agent is to give it up when faced with new information. In other words, $A \leq B$ holds just in case the agent's determination to keep B in the face of new information is at least as strong as his determination to keep A. The axioms (≤1) - (≤5) should be viewed as rationality requirements on such a relation for logically omniscient agents.

Clearly, the relation ≤ should be transitive, that is, we have (≤1). Furthermore, if $A \vdash_L B$, then the agent being logically omniscient knows that he cannot keep A without keeping B and he also knows that he cannot give up B without giving up A. Being rational he realizes that he cannot fight what is unavoidable. Thus, he should be at least as willing to give up A as to give up B (alternatively, he should be at least as determined to keep B as to keep A). This line of reasoning yields Axiom (≤2).

Axiom (≤3) says that if the agent's determination to keep B is at least as strong as his determination to keep A and his determination to keep C is also at least as strong as his determination to keep A, then the agent's determination to keep the conjunction $B \wedge C$ is at least as strong as his determination to keep A. For, clearly, if he were to give up $B \wedge C$, he would give up either B or C (or both), but his preparedness to give up A is at least as strong.

Intuitively, the relation of epistemic entrenchment is only defined for sentences that are accepted, i.e., that belong to G. It is, however, convenient to extend it to sentences that are not in G by stipulating that those sentences are less entrenched than all the sentences in G. This convention is expressed in Axiom (≤4). In the presence of (≤1) and (≤2), Axiom (≤4) is equivalent to the conjunction of the following two conditions:

(≤4.1) If $A \notin G$ and $B \in G$, then $A < B$.

(≤4.2) If $A \notin G$ and $B \notin G$, then $A \approx B$.

It follows from (≤2) that if $\vdash_L$ A, then B ≤ A, for all B. In particular, we have T ≤ A in this case. Axiom (≤5) says that the logical truths are the only sentences A such that T ≤ A. In other words, the logical truths are precisely those sentences that are maximally entrenched. Intuitively, those are only the sentences that the agent is determined to keep under any circumstances.

Observe that we differ from Gärdenfors (1988) in not assuming epistemic entrenchment relations to be connected. If we were to add:

A ≤ B or B ≤ A *(Connectedness)*

our axioms would be equivalent to Gärdenfors' postulates (EE1) - (EE5). In particular, connectedness together with (≤1) and (≤3) would imply Gärdenfors' postulate:

(EE3) A ≤ A ∧ B or B ≤ A ∧ B *(Conjunctiveness)*

Proof: By connectedness, we have: A ≤ B or B ≤ A. Suppose now that A ≤ B (the other case is similar). Connectedness also yields A ≤ A. Hence, by (≤3), A ≤ A ∧ B.

□

Conversely, (EE3) together with Transitivity and Dominance, which are among Gärdenfors' postulates, imply Connectedness.

(EE3) and Dominance yield:

A ≈ A ∧ B or B ≈ A ∧ B,

that is, A ∧ B is equally entrenched as A or equally entrenched as B. In fact, on Gärdenfors' approach, A ∧ B is equally entrenched as the least entrenched sentences among A and B. On our approach, however, A and B may be incomparable with respect to epistemic entrenchment, in which case A ∧ B is strictly less entrenched than both A and B. Instead of Gärdenfors' (EE3), we have:

if A ≤ B, then A ≤ A ∧ B.

In fact, we have the following "definition" of ≤ in terms of ≈:

A ≤ B iff A ≈ A ∧ B.

In the following, we let G be any fixed L-theory and ≤ an epistemic entrenchment ordering for G.

Definition 3.2.

(a) A set Γ of sentences is a *filter* (relative to ≤) if it satisfies the following conditions:

(i) Γ ≠ ∅;

(ii) If A ∈ Γ and A ≤ B, then B ∈ Γ;

(iii) If A, B $\in \Gamma$, then A $\wedge$ B $\in \Gamma$.

Alternatively, we can characterize a filter as a non-empty set Γ of sentences which satisfies the following two conditions:

(iv) A, B $\in \Gamma$ iff A $\wedge$ B $\in \Gamma$;

(v) If A $\approx$ B, then A $\in \Gamma$ iff B $\in \Gamma$.

It is easily seen that these characterizations are equivalent. We say that a filter Γ is *proper* if $\bot \notin \Gamma$.

(b) We let **F**($\leq$) be the set of all filters. We also refer to the elements of **F**($\leq$) as *fallbacks* (with respect to $\leq$). Proper filters are also called *proper* fallbacks.

(c) A family of sets $\mathcal{X}$ is said to be *directed* (with respect to $\subseteq$) if for any Γ and Δ in $\mathcal{X}$, there is a Σ in $\mathcal{X}$ such that $\Gamma \subseteq \Sigma$ and $\Delta \subseteq \Sigma$. In other words, $\mathcal{X}$ is directed if every finite subset of $\mathcal{X}$ has an upper bound in $\mathcal{X}$.

Intuitively, a proper fallback (w. r. t. an entrenchment ordering $\leq$ for G) is a weakening of G — a subtheory of G — which we obtain by removing from G all sentences that are not "sufficiently entrenched". We choose as our entrenchment standard some non-empty subset Γ of G which is closed under conjunctions and whose elements are all sufficiently entrenched for our purposes and we keep only those sentences that are at least as entrenched as some sentences in Γ.

Lemma 3.3. The set **F**($\leq$) of all filters (or fallbacks) with respect to $\leq$ has the following properties:

(**F**1) Every element of **F**($\leq$) is an L-theory.

(**F**2) If $\mathcal{X}$ is a subset of **F**($\leq$), then $\cap\mathcal{X} \in$ **F**($\leq$).

(**F**3) If $\mathcal{X}$ is a non-empty directed set of elements in **F**($\leq$), then $\cup\mathcal{X} \in$ **F**($\leq$).

(**F**4) If H $\in$ **F**($\leq$) and $\bot \notin$ H, then H $\subseteq$ G.

(**F**5) L $\in$ **F**($\leq$).

(**F**6) G $\in$ **F**($\leq$).

For the proof of this lemma and the proofs of the main lemmas and theorems to follow, see the Appendix.

(**F**2) says that **F**($\leq$) is a *closure system.*[4] A closure system satisfying condition (**F**3) is called *algebraic.* **F**($\leq$) is a complete lattice (with respect to the ordering under inclusion), where for each $\mathcal{X} \subseteq$ **F**($\leq$), inf($\mathcal{X}$) = $\cap\mathcal{X}$ and sup($\mathcal{X}$) = $\bigcap\{H \in$ **F**($\leq$): $\cup\mathcal{X} \subseteq H\}$. The *bottom* (or least element) of **F**($\leq$) is L and its *top* (or greatest element) is the inconsistent L-theory $K_\bot$ (= $\cap\varnothing$). If G is consistent, then the proper or consistent elements of **F**($\leq$) are exactly those that are included in G. In this case, the set of all proper filters is directed and G is its union. It should also be noted that if $\mathcal{X}$ is a non-empty directed set of *proper* filters, then $\cup\mathcal{X}$ is a proper filter.

For any set of sentences Γ, there exists a smallest member of **F**(≤) which includes Γ, namely

$$\uparrow(\Gamma) = \bigcap\{H \in \mathbf{F}(\leq): \Gamma \subseteq H\}.$$

We refer to ↑(Γ) as *the filter generated by* Γ.

Lemma 3.4. If $\Gamma \neq \emptyset$, then

$$\uparrow(\Gamma) = \{A: \exists\, B_1,\ldots, B_n \in \Gamma\ (n \geq 1),\ B_1 \wedge \ldots \wedge B_n \leq A\}.$$

Lemma 3.5. The operation $\uparrow: \mathbf{P}(\Phi) \rightarrow \mathbf{P}(\Phi)$ satisfies the following conditions:

(↑1) $\Gamma \subseteq \uparrow(\Gamma)$.
(↑2) $\uparrow(\Gamma) = \uparrow(\uparrow(\Gamma))$.
(↑3) If $\Gamma \subseteq \Delta$, then $\uparrow(\Gamma) \subseteq \uparrow(\Delta)$.
(↑4) $Cn_L(\Gamma) \subseteq \uparrow(\Gamma)$.
(↑5) If $\bot \notin G$, then $\bot \notin \uparrow(\Gamma)$ iff $\Gamma \subseteq G$.
(↑6) $\uparrow(G) = G$.
(↑7) $\uparrow(\emptyset) = L$.
(↑8) $\uparrow(\{A \wedge B\}) = \uparrow(\{A, B\})$.
(↑9) If $\Gamma \neq \emptyset$, then $\uparrow(\Gamma) = \bigcup\{\uparrow(\Delta): \Delta \subseteq \Gamma$ and Δ is finite$\}$.

Moreover, $\mathbf{F}(\leq) = \{\Gamma \subseteq \Phi: \uparrow(\Gamma) = \Gamma\}$.

The proof of Lemma 3.5 is omitted here. (The proofs of (↑8) and (↑9) use Lemma 3.4.)

Lemma 3.6. $A \leq B$ iff for every $H \in \mathbf{F}(\leq)$, if $A \in H$, then $B \in H$.

Proof: The left-to-right direction is trivial. To prove the other direction, assume that not($A \leq B$). Consider $F = \{C: A \leq C\}$. We want to show that F is a filter such that $A \in F$ and $B \notin F$. $A \in F$ by (≤2), and $B \notin F$ by the assumption. From (≤1) follows that F is closed upwards under ≤. Let $C, D \in F$. Then $A \leq C$ and $A \leq D$. So by axiom (≤3), $A \leq C \wedge D$, which means that $C \wedge D \in F$. Hence, F is a filter □

Corollary 3.7. $A \leq B$ iff $\uparrow(\{B\}) \subseteq \uparrow(\{A\})$.

Definition 3.8. Let G be an L-theory. A *fallback family* for G is a family **F** of sets of sentences satisfying the conditions (**F**1) - (**F**6) of Lemma 3.3.

Let G be a given L-theory. For any epistemic entrenchment ordering ≤ for G, we call **F**(≤) *the fallback family for G determined by* ≤, where **F**(≤) is the family of all filters with respect to ≤. As we have seen above, Lemma 3.3, **F**(≤) is indeed a fallback family. We also saw (Lemma 3.6) that the relation ≤ is definable in terms of **F**(≤). Thus, every epistemic entrenchment ordering ≤ for G can be thought of as being defined from a fallback family, namely from the family **F**(≤). Suppose now that ≤ and ≤′ are epistemic entrenchment relations such that **F**(≤) = **F**(≤′). Via Lemma 3.6 we then get that ≤ = ≤′. That is, the relation of determination between entrenchment relations and fallback families is one-to-one.

Next, we want to show that for *any* fallback family **F**, the relation ≤ defined by:

(D) A ≤ B iff for all H ∈ **F** if A ∈ H, then B ∈ H.

is an epistemic entrenchment relation and that, furthermore, **F** = **F**(≤). That is, there is a one-to-one correspondence between epistemic entrenchment relations and fallback families (for a given G).such that for any entrenchment relation the corresponding fallback family is set-theoretically definable from it and vice versa.

Lemma 3.9. Let **F** be any fallback family for G and ≤ the corresponding ordering of the sentences in $\mathcal{L}$ defined by condition (D). Then ≤ is an epistemic entrenchment ordering for G.

Proof: We prove that ≤ satisfies the postulates (≤1) - (≤5) for entrenchment orderings. The easy verifications of (≤1) - (≤3) using (**F**1) and (D) are omitted.

(≤4) Suppose ⊥ ∉ G. Assume also that A ∉ G. Consider any H ∈ **F** such that A ∈ H. Then not(H ⊆ G), so by (**F**4), ⊥ ∈ H. Hence, by (D), A ≤ ⊥. For the other direction, assume that A ≤ ⊥. If A ∈ G, then by (**F**6) and (D), ⊥ ∈ G. Hence, A ∉ G.

(≤5) Suppose T ≤ A. (**F**5) and (D) then yield A ∈ L. □

In view of the previous Lemma, for any fallback family **F** for G, we may speak of the ordering ≤ which is defined from **F** via (D) as *the epistemic entrenchment ordering corresponding to* **F**.

Lemma 3.10. Let **F** be a fallback family for G and let ≤ be the corresponding entrenchment ordering. Then **F** = **F**(≤). That is, the elements of **F** are precisely the filters with respect to the corresponding entrenchment ordering.

Proof: We first prove that **F** ⊆ **F**(≤). Let H ∈ **F**. Using (**F**1) and (D), it is easy to verify that H is a filter with respect to ≤. That is, H ∈ **F**(≤).

For the other direction, let H ∈ **F**(≤). Then using Lemma 3.5, we get

$$H = \bigcup\{\uparrow(\Delta): \Delta \subseteq H \text{ and } \Delta \text{ is finite}\} = \bigcup\{\uparrow(\{A\}): A \in H\}$$

However, for any A, $\uparrow(\{A\}) = \{B: A \leq B\}$, by Lemma 3.4. Hence, by (D), for any B, B $\in \uparrow(\{A\})$ iff for all K $\in$ **F**, if A $\in$ K, then B $\in$ K. That is, $\uparrow(\{A\}) = \bigcap\{K \in \mathbf{F}: A \in K\}$. (**F**2) then yields that $\uparrow(\{A\}) \in \mathbf{F}$, for any A. Now, $\{\uparrow(\{A\}): A \in H\}$ is a non-empty directed set. For let $A_1, A_2 \in H$. Then $\uparrow(\{A_1\}), \uparrow(\{A_2\}) \subseteq \uparrow(\{A_1 \wedge A_2\})$. Hence, by (**F**3), $\bigcup\{\uparrow(\{A\}): A \in H\} \in \mathbf{F}$. That is, H $\in$ **F**. □

Lemma 3.11. An epistemic entrenchment ordering ≤ is connected iff the corresponding fallback family **F**(≤) is a chain (i.e., is linearly ordered by inclusion).

Proof: Suppose the entrenchment ordering ≤ for G is connected, i.e., for all sentences A, B

(∗) $\quad$ A ≤ B or B ≤ A.

We want to show that for all H, K $\in$ **F**(≤), H $\subseteq$ K or K $\subseteq$ H. Suppose that neither H $\subseteq$ K nor K $\subseteq$ H. Then there are A, B such that A $\in$ H, A $\notin$ K, B $\in$ K and B $\notin$ H. But then by the definition of a filter, neither A ≤ B nor B ≤ A, contrary to (∗).

For the other direction, assume that **F**(≤) is linearly ordered by inclusion. Consider any sentences A, B. Then $\uparrow(\{B\}) \subseteq \uparrow(\{A\})$ or $\uparrow(\{A\}) \subseteq \uparrow(\{B\})$. This implies, by Corollary 3.7, that A ≤ B or B ≤ A. □

Notice that if **F**(≤) is a chain, then every non-empty subset of **F**(≤) is directed. It follows that **F**(≤) in this case is closed under the formation of arbitrary unions of non-empty subsets.

Definition 3.12. Let ≤ be an epistemic entrenchment ordering for G. A fallback H in **F**(≤) is said to be *A-permitting* if ¬A $\notin$ H. H $\in$ **F**(≤) is a *maximal A-permitting fallback (for G)* if:

(i) $\quad$ ¬A $\notin$ H; and

(ii) $\quad$ for any K $\in$ **F**(≤), if H $\subset$ K, then ¬A $\in$ K.

Lemma 3.13. Let ≤ be an epistemic entrenchment ordering for G. Every A-permitting fallback in **F**(≤) is included in a maximal A-permitting fallback.

Proof: Let H $\in$ **F**(≤) such that ¬A $\notin$ H. We let

$$\mathcal{K} = \{K \in \mathbf{F}(\leq): H \subseteq K \text{ and } \neg A \notin K\}.$$

$\mathcal{K} \neq \varnothing$, since H $\in \mathcal{K}$. Let $\varnothing \neq \mathcal{C} \subseteq$ K be a chain in $\mathcal{K}$. Then $\mathcal{C}$ is a directed set, so by (**F**3), $\cup\mathcal{C} \in \mathbf{F}(\leq)$. Clearly, H $\subseteq \cup\mathcal{C}$. Moreover, ¬A $\notin \cup\mathcal{C}$, for otherwise, for some K $\in \mathcal{C}$, ¬A $\in$ K, which is impossible. Hence, $\cup\mathcal{C} \in \mathcal{K}$. We have shown that every non-empty chain in $\mathcal{K}$ has an upper bound in $\mathcal{K}$. Zorn's lemma then yields that $\mathcal{K}$ has a maximal element. □

We shall end this section by showing that entrenchment orderings that do not satisfy connectedness can be represented by *families* of connected entrenchment orderings. The basic idea, due to Hans Rott (personal communication), is the following: Consider any non-empty set X_G of entrenchment orderings for G. Then, as is easily seen, $\cap X_G$ is an entrenchment ordering for G. That is, $\cap X_G$ satisfies the postulates (≤1) - (≤5). Now, if X_G is a non-empty family of *connected* entrenchment orderings for G, then, of course, $\cap X_G$ is an entrenchment ordering for G. However, $\cap X_G$ need not be connected!

This idea gives rise to a question: Can any epistemic entrenchment ordering ≤ for G be represented as $\cap X_G$ for some non-empty family X_G of *connected* entrenchment orderings for G? Below, we show that the answer to this question is affirmative.

Consider any entrenchment ordering ≤ for G. We say that ≤′ is a *connected extension* of ≤ (relative to G) if ≤′ satisfies the following conditions:

(i) $\leq \subseteq \leq'$;

(ii) ≤′ is a connected entrenchment ordering for G.

≤′ is a *minimal connected extension* of ≤, if it in addition to (i) and (ii) satisfies:

(iii) for any ≤″, if ≤″ satisfies (i) and (ii), then $\leq'' \not\subset \leq'$.

We let E(≤) and ME(≤) be the set of all connected extensions and the set of all minimal connected extensions of ≤, respectively.

The next theorem says that an entrenchment ordering ≤ for G that is not connected can be represented by any set X of connected extensions of ≤ that contains all the minimal connected extensions.

Theorem 3.14. If ≤ be an entrenchment ordering for G and let X be any set of connected entrenchment orderings such that $ME(\leq) \subseteq X \subseteq E(\leq)$. Then $\leq = \cap X$.

According to the above theorem, there are in general many different sets of connected entrenchment orderings for G that have the same non-connected ordering as their intersections. Therefore, the representation of epistemic entrenchment in terms of sets of connected relations is more *finegrained* or *discriminating* than the representation in terms of non-connected relations. But perhaps such a representation would be *too* discriminating. From the intuitive point of view, it is easy to understand what it means to say that the agent's entrenchment ordering is incomplete (non-connected): from his point of view, there are propositions that are *incomparable* with respect to entrenchment; perhaps because these propositions are so different from each other, or perhaps because they are totally unrelated. But what is the meaning of the claim that the agent's state consists of a *class* of entrenchment orderings? Is the idea that the agent vacillates between different entrenchment orderings? That he sometimes changes his beliefs according to one ordering and sometimes according to another? Or is it rather that his epistemic state is *indeterminate*? But then why not simply say that his entrenchment ordering is incomplete?

There is an analogy here with the probabilistic representation of an agent's epistemic state. The agent's belief state can be represented as (i) a definite real-valued probability function over propositions (the classical Bayesian approach), or (ii) a probability function that takes intervals as values, or as (iii) a class of real-valued probability functions. A class Π of ordinary probability functions, if it satisfies certain conditions, determines an interval-valued probability function. We assigne to a proposition A the interval [inf(P(A)), sup(P(A))], where P varies over the members of Π. It is clear, of course, that different classes may induce the same probability intervals. One could say that Gärdenfors' approach in terms of a single connected entrenchment ordering corresponds to alternative (i) above; our approach using an incomplete entrenchment ordering is like the alternative (ii); while working with a class of connected entrenchment orderings is similar to alternative (iii). But entrenchment relations are not like probability orderings: the axiom (≤3) — Conjunctive Closure — would be absurd if ≤ were interpreted as 'at most as probable as'. Therefore, the analogy should not be taken too far.

4. Relational Revision

The Alchourrón-Gärdenfors-Makinson approach treats belief revision as a function on theories: for any theory G and any sentence A, there is a unique revision of G with A. If we think of revision as being defined out of an epistemic entrenchment ordering ≤ of sentences, then the functional approach is a natural consequence of the assumption that epistemic entrenchment is connected. Intuitively, when the agent in belief state G receives a new piece of information A, he selects a maximal A-permitting fallback for G and adds A to it. That is, we have the following definition of revision in terms of epistemic entrenchment.

Definition 4.1. Let ≤ be an epistemic entrenchment ordering for G. We say that H is a *revision* of G with A relative to ≤ if either ¬A ∈ L and H = Φ or there exists some maximal A-permitting fallback K in **F**(≤) such that H = K + A. (Here, K + A, the *expansion* of K with A, is $Cn_L(K \cup \{A\})$.)

If ≤ is connected, then the family **F**(≤) of all fallbacks for G is a chain, so there is a *unique* maximal A-permitting fallback for G. In fact, this fallback is the set of all B such that $B \nleq \neg A$. Accordingly, in this case there is a unique revision of G with A. If, however, we give up the requirement of connectedness, that is, if we allow sentences to be *incomparable* with respect to epistemic entrenchment, then there may be several maximal A-permitting fallbacks and hence also several revisions of G with A. Revision becomes relational rather than functional.

In the next definition, we formulate a system of axioms for relational belief revision.

Definition 4.2. A *(relational) belief revision system (b.r.s.)* is an ordered pair <L, **R**> such that L is a consistent logic, $\mathbf{R} \subseteq \mathcal{T}_L \times \Phi \times \mathcal{T}_L$ and the following requirements are satisfied for all L-theories G, H and all sentences A, B:

(**R**1) $(\exists H \in \mathcal{T}_L)(G\mathbf{R}_A H)$. *(Seriality)*

(**R**2) If $G\mathbf{R}_A H$, then $A \in H$. *(Success)*

(**R**3) If $\neg A \notin G$ and $G\mathbf{R}_A H$, then $H = G + A$. *(Expansion)*

(**R**4) If $\neg A \notin L$ and $G\mathbf{R}_A H$, then $\bot \notin H$. *(Strong Consistency)*

(**R**5) If $\vdash_L A \leftrightarrow B$, then $G\mathbf{R}_A H$ iff $G\mathbf{R}_B H$. *(Substitutivity)*

(**R**6) If $G\mathbf{R}_A H$ and $\neg B \notin H$, then $G\mathbf{R}_{A\wedge B}(H + B)$ *(Revision by Conjunction)*

(**R**7) If $G\mathbf{R}_A H$ and $\forall K(G\mathbf{R}_{A\vee B}K \rightarrow \neg A \notin K)$, then $\exists K(G\mathbf{R}_{A\vee B}K$ and $H = K + A)$.

We read $G\mathbf{R}_A H$ as "H is a (possible) revision of G with A".

In the presence of *functionality:*

If $G\mathbf{R}_A H$ and $G\mathbf{R}_A K$, then K = H,

the axioms (**R**1) - (**R**7) are equivalent to Gärdenfors' (1988) axioms (K*1) - (K*8) for belief revision.

Lemma 4.3. If <L, **R**> is a belief revision system, then it satisfies the following principles:

(**R**8) If $G\mathbf{R}_{A\wedge B}H$ and $\forall K(G\mathbf{R}_A K \rightarrow \neg B \notin K)$, then $\exists K(G\mathbf{R}_A K$ and $H = K + B)$

(**R**9) If (i) $G\mathbf{R}_{A\vee B}H$ and $\neg B \in H$; and (ii) $\forall K(G\mathbf{R}_{B\vee C}K \rightarrow \neg B \notin K)$, then $G\mathbf{R}_{A\vee C}H$ and $\neg C \in H$. *(Transitivity)*

Moreover, in the presence of the other axioms (in fact, only (**R**2) and (**R**5) are needed), (**R**7) and (**R**8) are equivalent principles.

As we shall see next, if belief revision is defined from a system of entrenchment orderings — one for each theory G — then it satisfies the axioms (**R**1) - (**R**7). Moreover, the original system of entrenchment orderings will be seen to be definable from the corresponding belief revision system. First, however, we introduce the notion of an entrenchment system and the equivalent notion of a fallback system.

Definition 4.4.

(a) An *(epistemic) entrenchment system (e.e.s.)* is an ordered pair $<L, \leq>$ such that L is a consistent logic and $\leq$ is a function which assigns to every L-theory G an epistemic entrenchment relation $\leq_G$ for G.

(b) A *fallback system* is an ordered pair $<L, \mathbf{F}>$ such that L is a consistent logic and **F** is a function which assigns to every L-theory G a fallback family $\mathbf{F}_G$ for G.

In view of the results in the previous section, there is a natural one-to-one correspondence between entrenchment systems and fallback systems. The next theorem states that every entrenchment system (or equivalently, every fallback system) determines a belief revision system.

Theorem 4.5. Let $<L, \leq>$ be an entrenchment system. Define $<L, \mathbf{R}>$ as follows: for all L-theories G, H and all sentences A,

($D_\mathbf{R}$) $\quad G\mathbf{R}_A H$ iff H is a revision of G with A relative to $\leq_G$.

Then $<L, \mathbf{R}>$ is a belief revision system, that is, it satisfies the axioms (**R**1) - (**R**7). In addition, the following equivalence holds for all A, B and G:

($D_\leq$) $\quad A \leq_G B$ iff either $B \in L$ or $\forall H$(if $G\mathbf{R}_{\neg(A\wedge B)}H$, then $A \notin H$).

That is, if the belief revision system $<L, \mathbf{R}>$ is defined from an entrenchment system $<L, \leq>$ by means of ($D_\mathbf{R}$), then $<L, \leq>$ can be recovered from $<L, \mathbf{R}>$ via ($D_\leq$).

Theorem 4.6. Let $<L, \mathbf{R}>$ be a belief revision system and define for every L-theory G and all A, B:

($D_\leq$) $\quad A \leq_G B$ iff either $B \in L$ or $\forall H$(if $G\mathbf{R}_{\neg(A\wedge B)}H$, then $A \notin H$).

The so constructed system $<L, \leq>$ is an epistemic entrenchment system.

Suppose that we consider a belief revision system <L, **R**>. Are our axioms on belief revision (cf. Definition 4.2) sufficient to guarantee that <L, **R**> is *representable,* that is, that there is some epistemic entrenchment system <L, ≤ > such that for all L-theories G, H and all sentences A,

($D_{\mathbf{R}}$) $G\mathbf{R}_A H$ iff H is a revision of G with A relative to $\leq_G$.

By Theorem 4.6, we know that in terms of <L **R**> we can define an entrenchment system <L, $\leq^{\mathbf{R}}$>. Suppose now that we define a belief revision system <L, **S**> from <L, $\leq^{\mathbf{R}}$>, where $G\mathbf{S}_A H$ holds just in case H is a revision of G with A relative to $\leq^{\mathbf{R}}_G$.

In view of Theorem 4.5, we know that R is representable if and only if **R** = **S**.

It can be shown that this representability problem has a negative answer. Not all belief revision systems are representable. The reason for this is that our set of axioms for **R** is not sufficiently strong to exclude an unintended interpretation of belief revision — an interpretation in which belief revisions are not required to be "minimal".

Thus, suppose that we have an entrenchment ordering ≤ for G and define the corresponding family **F**(≤) of fallbacks. Now, suppose that we define belief revision in terms of **F**(≤), not as we have done before (cf. Definition ($D_{\mathbf{R}}$) in Theorem 4.5), but as follows:

(D'_A) $G\mathbf{R}'_A$ H iff either

(i) $\neg A \in$ L and $\bot \in$ H; or

(ii) $\neg A \notin$ G and H = G + A; or

(iii) $\neg A \in$ G and there is a (not necessarily maximal) A-permitting fallback K in **F**(≤) such that H = K + A.

It is easy to check that all the belief revision axioms (**R**1) - (**R**7) hold for **R**′ given this interpretation. At the same time, the *intended* relation **R** is included in the non-intended **R**′.

In general, **R** will be *properly* included in **R**′. This can be seen if one starts from a connected entrenchment ordering ≤. Then **R** but *not* **R**′ will be functional.

Now, we are going to show that **R**′ will not in general be representable. Suppose that **R** ⊂ **R**′ and that **R**′ is representable. Then there exists an entrenchment ordering ≤′ such that **R**′ is definable from ≤′ via definition ($D_{\mathbf{R}}$). But then by Theorem 4.5, ≤′ is definable from **R**′ by means of ($D_{\leq}$). We are now going to show that ≤ = ≤′. This will entail a contradiction, because the revision relation **R** obtainable from ≤ via ($D_{\mathbf{R}}$) differs from the revision relation **R**′ obtainable from ≤′ via the same definition.

<u>Proof that if **R**′ is representable, then ≤ = ≤′</u>.

By ($D_{\leq}$) and the fact that $\mathbf{R} \subseteq \mathbf{R}'$, it follows immediately that $\leq' \subseteq \leq$. For the other direction, suppose that $A \leq B$. In the limiting case when $B \in L$, it is immediate that $A \leq' B$. Thus, in order to prove $A \leq' B$, for the principal case when $B \notin L$, suppose that $G\mathbf{R}'_{\neg A \vee \neg B} H$. We have to show that $A \notin H$. Suppose that $A \in H$. Then there must exist a $(\neg A \vee \neg B)$-permitting fallback K in $\mathbf{F}(\leq)$ such that $H = K + (\neg A \vee \neg B)$. But then by lemma 3.13, there exists a maximal $(\neg A \vee \neg B)$-permitting fallback T such that $K \subseteq T$. Since, $A \in H$, (i) $A \in T + (\neg A \vee \neg B)$. But, by the definition of $\mathbf{R}$, we have that (ii) $G\mathbf{R}_{\neg A \vee \neg B}(T + (\neg A \vee \neg B))$. But, (i) and (ii) contradict the assumption that $A \leq B$.

□

We have seen that in general belief revision systems are not representable. This negative result naturally gives rise to the question: how can our system of axioms for relational belief revision be strengthened (in a nice way) so as to guarantee representability?

The situation is different, however, if we want to add functionality: functional belief revision systems are always representable. The theorem below is essentially equivalent to Grove's (1988) representation theorem for functional belief revision in terms of families of spheres. Our proof, however, is different from his. (See the Appendix.)

Theorem 4.7. Let <L, **R**> be a functional belief revision system. Define the corresponding entrenchment system <L, ≤ > by means of ($D_{\leq}$) and let <L, **S**> be the belief revision system obtained from <L, ≤ > via ($D_{\mathbf{R}}$). Then $\mathbf{R} = \mathbf{S}$.

At the end of the previous section, we have discussed the possibility of representing the agent's belief state as a set X_G of connected entrenchment orderings for a given belief set G. If one were to choose this approach, one would have two alternative ways of defining relational belief revision.

Alternative 1. Go to the intersection $\cap X_G$. As we already know, $\cap X_G$ is an entrenchment ordering. Then define **R** in the standard way from $\cap X_G$ via $\mathbf{F}(\cap X_G)$.

Alternative 2. For any $\leq'$ in X_G, define the corresponding revision relation $\mathbf{R}^{\leq'}$ via $\mathbf{F}(\leq')$. Since each such $\leq'$ is connected, $\mathbf{R}^{\leq'}$ will of course be a function. Then, we define:

$$G\mathbf{S}_A H \text{ iff for some } \leq' \text{ in } X_G,\ G\mathbf{R}^{\leq'}_A H.$$

It is easy to check that **S** satisfies all our axioms for relational belief revision.

Alternative 2 is more liberal than alternative 1. We get more possible revisions that way. In order to see this, the reader should note that every fallback in $\mathbf{F}(\cap X_G)$

is a fallback in **F**(≤′), for some ≤′ in X_G. Therefore, every maximal A-permitting fallback in **F**($\cap X_G$) is a maximal A-permitting fallback in **F**(≤′), for some ≤′ in X_G. But the converse need not hold: a maximal A-permitting fallback in **F**(≤′) may not be a *maximal* A-permitting fallback in **F**($\cap X_G$).

In what follows, we keep to our preferred interpretation of relational belief revision in terms of a single epistemic entrenchment ordering.

5. Fallbacks and Contractions

According to Isaac Levi (1980, sect. 3.5), belief revision should be seen as a two-step process. In order to revise G with A, we first *contract* G with ¬A and then expand the resulting belief set with A. On our approach, belief revision is also a two-step process. A possible revision of G with A is obtained by first moving to a maximal A-permitting fallback of G and then expanding this fallback with A. This similarity between the two descriptions of belief revision might suggest the identification of contractions with maximal fallbacks: one might think that H is a possible contraction of G with A just in case H is a maximal ¬A-permitting fallback of G. (If A is a logically true, we let G itself be the only contraction of G with A.) This tentative suggestion appears in Rott (1989).[5]

If we look at theories as sets of possible worlds (where the latter might be identified with L-maximal theories), then the following figure illustrates the present proposal:

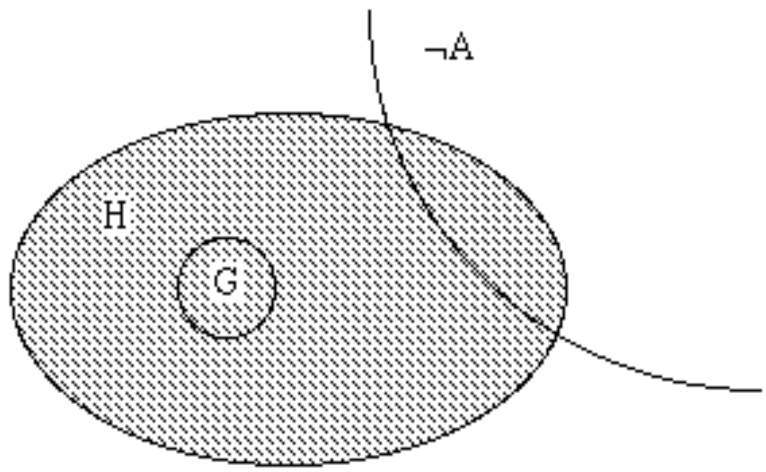

Figure 4.

This way of defining contractions might invite an objection: According to the Alchourrón-Gärdenfors-Makinson approach, contraction is assumed to satisfy the postulate of *recovery*. In terms of our relational approach, this postulate amounts to the following claim:

(Recovery) If $A \in G$ and $G\mathbf{R}_A^- H$, then $H + A = G$,

where $G\mathbf{R}_A^- H$ should be read as: "H is a possible contraction of G with A".

According to Recovery, if one contracts a theory with A and then expands the contraction with A, one recovers the theory one has started with.
Now, it is easy to see that if one moves from a theory G containing A to a maximal ¬A-permitting fallback H and then expands H with A, one will not normally recover G.[6] In order to see this, let A and B be two equally entrenched sentences in G such that their disjunction is logically true, but they are not contradictories (For instance, let A and B be "X is good or neutral" and "X is neutral or bad", respectively.) Since A and B are equally entrenched, it follows from the definition of fallbacks in terms of entrenchment that any ¬A-permitting fallback H must be ¬B-permitting as well. Now, consider $H + A$. It is easy to see that $B \notin H + A$. Otherwise, we would have that $A \to B \in H$, which in view of the fact that $A \vee B \in L$ and that $B \notin H$, is impossible. But since $B \in G$, it follows that $G \neq H + A$. To illustrate:

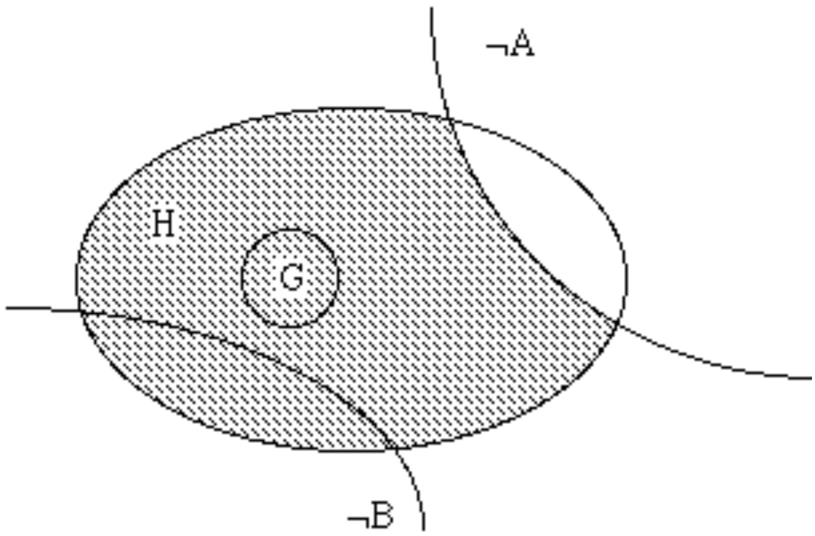

Figure 5.

It is obvious that in this example, $H + A$ differs from G.

Thus, if belief contraction should satisfy Recovery, contractions cannot be identified with maximal fallbacks.

However, the strength of this objection may well be doubted. Recovery is not a particularly convincing condition on contraction. As is easily seen, it has the following consequence:

If $B \in G$ and $B \vdash_L A$, then $G\mathbf{R}^-_A H$ implies that $B \in H + A$.

This principle is quite surprising. According to it, if we believe that Jesus was God's son and then retract our belief in Jesus' historical existence, then we are going to recover our belief Jesus' divine origin upon learning that Jesus actually existed. It is difficult to see why this should be the case. For a sustained criticism of Recovery from this point of view, see Niederée (1990). It should also be noted that Isaac Levi has never accepted the Recovery Postulate for contraction.

However, there is another objection against identifying contractions with maximal fallbacks that we find much more convincing. As we know, fallbacks have the following property: if a fallback H is ¬A-permitting and B is at most as entrenched as A, then H is ¬B-permitting as well (This is an easy consequence of the fact that fallbacks are closed upwards under entrenchment.) Therefore, if we were to identify contractions with maximal fallbacks, we would have to accept the following principle:

If $G\mathbf{R}^-_A H$ and $B \leq_G A$, then $B \notin H$.

That is, when contracting with A, we would have to give up all the beliefs that are at most as entrenched as A, even those that are quite unrelated to our original belief in A. For example if we were called upon to give up some very entrenched belief of ours, for instance that the Earth rotates around the Sun, we would have to give up all beliefs that are less entrenched, even though they might be quite independent of the belief that is being contracted. For example, we would no longer be able to believe that Salman Rushdie has written "Satanic Verses".

This is not how we do our contractions. Thus, the identification of contractions with maximal fallbacks does not seem to work. There is an alternative, however, which we should consider and which goes back to Grove (1988). We have used maximal fallbacks to define revisions. We might now also use them to define contractions in the following way:

(Grove) $G\mathbf{R}^-_A H$ iff $A \in L$ and $G = H$, or there is a maximal ¬A-permit- ting fallback K in $\mathbf{F}(\leq_G)$ such that $H = G \cap (K + \neg A)$.

In the figure below, the shaded area represents Grove's contraction. The theory $G \cap (K + \neg A)$ is represented by the union of the sets of worlds corresponding to G and $K + \neg A$.

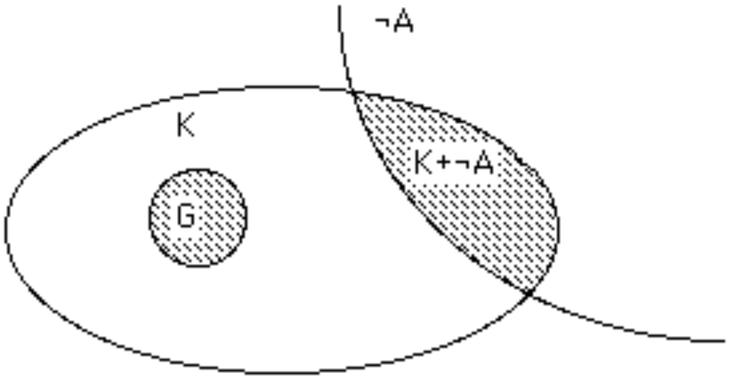

Figure 6.

Given our interpretation of revision in terms of fallbacks, *(Grove)* amounts to the following definition of contraction in terms of revision:

(Harper) $G\mathbf{R}_A^- H$ iff there is some K such that $G\mathbf{R}_{\neg A}K$ and H = $G \cap K$.

(Harper) is the relational version of the so-called *Harper Identity:* $G_A^- = G \cap G_{\neg A}^*$. (Cf. Gärdenfors, 1988, sect. 3.6; G_A^- and G_A^* here stand for functional contraction and functional revision of G with A, respectively.)

(Grove) also yields Levi's reduction of revision to contraction, namely

(Levi) $G\mathbf{R}_A H$ iff for some K, $G\mathbf{R}_{\neg A}^- K$ and $H = K + A$.

Unfortunately, Grove's proposal implies Recovery for contraction. In fact, Recovery follows from *Harper* (given Success), and as we have seen, *Harper* follows from *Grove*. It seems, therefore, that Grove's contractions are too strong from the intuitive point of view.

On the other hand, we have seen that our original proposal (contractions = maximal fallbacks) gives us contractions that are too weak: we loose all the propositions that are at most as entrenched as the proposition we contract with. One would like to say that the truth lies somewhere in between the two extremes: the original proposal and Grove's definition seem to give us the "lower" and the "upper" limit for contraction. This idea could be expressed as the following set of adequacy conditions on the contraction relation:

(C1) If $G\mathbf{R}_A^- H$ and $A \notin L$, then there exists some maximal ¬A-permitting fallback K in $\mathbf{F}(\leq_G)$ such that $K \subseteq H \subseteq G \cap (K + \neg A)$

(C2) For every maximal ¬A-permitting fallback K in $\mathbf{F}(\leq_G)$, there exists some L-theory H such that $G\mathbf{R}^-_A H$ and $K \subseteq H \subseteq G \cap (K + \neg A)$.

(C3) If $A \in L$, then $G\mathbf{R}^-_A H$ iff $G = H$.

(C3) takes care of the limiting case. The figure below illustrates the principal case, in which $A \notin L$:

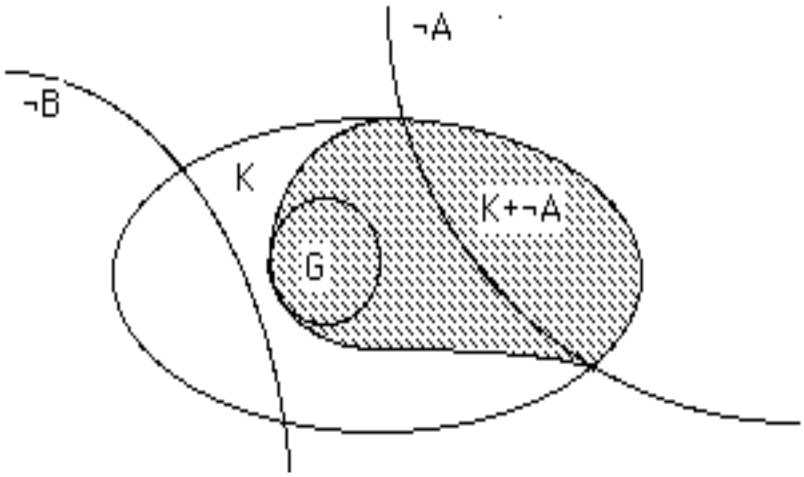

Figure 7.

As is easily seen, K and $G \cap K + \neg A$ are the lower and the upper limit for the contraction H represented by the shaded area.

Clearly, the present proposal does not imply Recovery (nor does it imply *Harper*, for that matter): in figure 7, H + A differs from the original theory G. Also, on this proposal, some propositions that are at most as entrenched as A may be retained in a contraction with A (cf. proposition B, in the figure).

Although it does not imply *Harper*, the present proposal entails the Levi identity.

Lemma 6.1. Let $<L, \leq>$ be an entrenchment system and let $\mathbf{R}$ be the corresponding revision relation defined from ≤ via (D≤). Suppose that $\mathbf{R}^- \subseteq \mathcal{T}_L \times \Phi \times \mathcal{T}_L$ is any relation satisfying (C1) - (C3). Then $\mathbf{R}$ and $\mathbf{R}^-$ satisfy *Levi.*

This fact shows that there is no fundamental incompatibility between our definition of revision from entrenchment and Levi's interpretation of revision in terms of contraction. Given that the contraction relation respects the entrenchment ordering in the way indicated by (C1) - (C3), the two interpretations of revision are going to coincide.

Appendix: Proofs of Lemmas and Theorems

Proof of Lemma 3.3: (**F**1) Let $\Gamma \in \mathbf{F}(\leq)$. Suppose $\Gamma \vdash_L A$. We have to show that $A \in \Gamma$. By the compactness of L, there are $B_1,..., B_n \in \Gamma$ such that $B_1,..., B_n \vdash_L A$. By the deduction theorem and propositional logic, $\vdash_L B_1 \wedge ... \wedge B_n \rightarrow A$. That is,.$B_1 \wedge ... \wedge B_n \vdash_L A$. (≤2) then yields that $B_1 \wedge ... \wedge B_n \leq A$. By the closure of Γ under ∧, $B_1 \wedge ... \wedge B_n \in \Gamma$. But then, since Γ is closed upwards under ≤, $A \in \Gamma$.

(**F**2) Let $\mathcal{X} \subseteq \mathbf{F}(\leq)$. Consider $\cap\mathcal{X}$. $L \subseteq \cap\mathcal{X}$, so $\cap\mathcal{X} \neq \emptyset$. Suppose that $A \in \cap\mathcal{X}$ and $A \leq B$. Then $A \in H$, for all H in $\mathcal{X}$. Consequently, $B \in H$, for all H in $\mathcal{X}$. That is, $B \in \cap\mathcal{X}$. We have shown that $\cap\mathcal{X}$ is closed upwards under ≤. Now, let $A, B \in \cap\mathcal{X}$. Then A and B both belong to all theories in $\mathcal{X}$. Hence, the same is true for $A \wedge B$. It follows that $A \wedge B \in \cap\mathcal{X}$. This completes the proof that $\cap\mathcal{X}$ is a filter.

(**F**3) Let $\mathcal{X}$ be a non-empty directed set of elements in $\mathbf{F}(\leq)$. Consider $\cup\mathcal{X}$. Clearly, $L \subseteq \cup\mathcal{X}$, so $\cup\mathcal{X} \neq \emptyset$. Suppose $A \in \cup\mathcal{X}$ and $A \leq B$. Then, for some $H \in \mathcal{X}$, $A \in H$. Since H is a filter, $B \in H$. It follows that $B \in \cup\mathcal{X}$. Now, let $A, B \in \cup\mathcal{X}$. Then there are H_1 and H_2 in $\mathcal{X}$ such that $A \in H_1$ and $B \in H_2$. Since $\mathcal{X}$ is directed, there is a $K \in \mathcal{X}$ such that $H_1 \subseteq K$ and $H_2 \subseteq K$. Hence, $A, B \in K$. But then $A \wedge B \in K$, since K is a filter. Hence, $A \wedge B \in \cup\mathcal{X}$.

(**F**4) Suppose H is a proper filter. We want to show that $H \subseteq G$. If G is inconsistent, there is nothing to prove. So, we assume that $\bot \notin G$. Suppose now that $A \in H$. Since H is proper, not($A \leq \bot$). It follows by (≤4) that $A \in G$. Hence, $H \subseteq G$.

(**F**5) Let $A \in L$ and $A \leq B$. Then $T \vdash_L A$. Using (≤2) we then get $T \leq A$. (≤1) yields $T \leq B$. This in turn implies $B \in L$ by means of (≤5). Since L is non-empty and closed under conjunction, it follows that L is a filter.

(**F**6) If G is inconsistent, then clearly $G \in \mathbf{F}(\leq)$. So assume that $\bot \notin G$. Suppose $A \in G$ and $A \leq B$. Using (≤4) we conclude that not($A \leq \bot$). This together with $A \leq B$ then implies not($B \leq \bot$) (using (≤1)). Applying (≤4) again, we get $B \in G$. Hence, G is closed upwards under ≤. Since, G is an L-theory, it is non-empty and closed under conjunctions. Hence, $G \in \mathbf{F}(\leq)$. □

Proof of Lemma 3.4: Let $\Delta = \{A: \exists B_1,..., B_n \in \Gamma\ (n \geq 1), B_1 \wedge ... \wedge B_n \leq A\}$. We want to show that: (i) $\Delta \subseteq \uparrow(\Gamma)$; and (ii) $\uparrow(\Gamma) \subseteq \Delta$.

(i) Let $H \in \mathbf{F}(\leq)$ such that $\Gamma \subseteq H$. Suppose that $\exists B_1,..., B_n \in \Gamma\ (n \geq 1)$, $B_1 \wedge ... \wedge B_n \leq A$. Then, by by the definition of a filter, $B_1 \wedge ... \wedge B_n \in H$. Hence, once again by the definition of a filter, $A \in H$. Thus, $\Delta \subseteq \uparrow(\Gamma)$.

(ii) It suffices to prove that Δ is a filter that includes Γ. Let $A \in \Gamma$. By (≤2), $A \leq A$, so $A \in \Delta$, by the definition of Δ. Thus $\Gamma \subseteq \Delta$. Since $\Gamma \neq \emptyset$, it follows that $\Delta \neq \emptyset$.

Suppose that $A \in \Delta$ and $A \leq B$. Then, for some $B_1,..., B_n \in \Gamma$, $B_1 \wedge ... \wedge B_n \leq A$. By (≤1) it follows that $B_1 \wedge ... \wedge B_n \leq B$. Hence, $B \in \Delta$.

Now, let $A_1, A_2 \in \Delta$. Then there are $B_1,..., B_n, C_1,..., C_m \in \Gamma$ such that $B_1 \wedge ... \wedge B_n \leq A_1$ and $C_1 \wedge ... \wedge C_m \leq A_2$. But $B_1 \wedge ... \wedge B_n \wedge C_1 \wedge ... \wedge C_m \leq B_1 \wedge ... \wedge B_n$ and $B_1 \wedge ... \wedge B_n \wedge C_1 \wedge ... \wedge C_m \leq C_1 \wedge ... \wedge C_m$ by (≤2). Hence, by (≤1), $B_1 \wedge ... \wedge B_n \wedge C_1 \wedge ... \wedge C_m \leq A_i$ $(i = 1, 2)$. It follows, using (≤3) that $B_1 \wedge ... \wedge B_n \wedge C_1 \wedge ... \wedge C_m \leq A_1 \wedge A_2$. Thus, $A_1 \wedge A_2 \in \Delta$. □

Proof of Theorem 3.14: Suppose that $ME(\leq) \subseteq X \subseteq E(\leq)$. That $\leq \subseteq \cap E(\leq)$ immediately follows from condition (i) in the definition of a minimal connected extension of ≤. But since $X \subseteq E(\leq)$, $\cap E(\leq) \subseteq \cap X$. Hence, $\leq \subseteq \cap X$.

In order to prove that $\cap X \subseteq \leq$, suppose that for some A and B, $A \nleq B$. We want to show that for some $\leq' \in ME(\leq)$, $A \nleq' B$. If $A \nleq B$, then for some fallback $K \in \mathbf{F}(\leq)$, $A \in K$ but $B \notin K$ (Cf. Lemma 3.6). We are now going to prove two claims:

Claim 1. If $K \in \mathbf{F}(\leq)$, then there is some maximal chain $\mathcal{C}$ in $\mathbf{F}(\leq)$ such that $K \in \mathcal{C}$.

Claim 2. If $\mathcal{C}$ is a maximal chain in $\mathbf{F}(\leq)$, then the entrenchment ordering $\leq'$ defined from $\mathcal{C}$ via:

(D) $A \leq' B$ iff for all $H \in \mathcal{C}$, if $A \in H$, then $B \in H$,

belongs to $ME(\leq)$.

Since we have assumed that $A \in K$ but $B \notin K$, it will follow from claims 1 and 2 that for some $\leq' \in ME(\leq)$, $A \nleq' B$. But then, since $ME(\leq) \subseteq X$, there is some $\leq' \in X$ such that $A \nleq' B$. It follows that $\leq \subseteq \cap X$.

To prove Claim 1, consider the set $\mathbf{K} = \{\mathcal{C}: \mathcal{C}$ is a chain in $\mathbf{F}(\leq)$ and $K \in \mathcal{C}\}$. It is easy to see that $\mathbf{K}$ is non-empty and that every chain $\mathbf{C} \subseteq \mathbf{K}$ has an upper bound in $\mathbf{K}$ (namely, $\cap\mathbf{C}$). Hence, by Zorn's lemma, there is a maximal element $\mathcal{C}$ in $\mathbf{K}$. Q.E.D.

To prove claim 2, consider a maximal chain $\mathcal{C}$ in $\mathbf{F}(\leq)$ and define $\leq'$ via (D). It is obvious that $\mathcal{C}$ is a fallback family for G. Therefore, by Lemma 3.9, $\leq'$ is an entrenchment ordering for G. Moreover, $\leq'$ is connected (cf. Lemma 3.11). It is also clear that $\leq \subseteq \leq'$. It remains to show that $\leq'$ satisfies condition (iii) in the definition of a minimal connected extension of ≤. Suppose it does not. Then, there is some connected entrenchment ordering $\leq''$ for G such that $\leq \subseteq \leq'' \subset \leq'$. Then, by lemma 3.11, $\mathbf{F}(\leq'')$ is a chain. Since $\leq \subseteq \leq''$, $\mathbf{F}(\leq'') \subseteq \mathbf{F}(\leq)$. Since $\leq'' \subset \leq'$, $\mathcal{C}$ $(= \mathbf{F}(\leq')) \subset \mathbf{F}(\leq')$. But this is impossible, since a maximal chain in $\mathbf{F}(\leq)$ cannot be properly included in another chain in $\mathbf{F}(\leq)$. □

Proof of Lemma 4.3: We first show how to obtain (**R**8) from (**R**7). In (**R**7) replace A by $A \wedge B$ and B by A. We then get:

If $G\mathbf{R}_{A\wedge B}H$ and $\forall K(G\mathbf{R}_{(A\wedge B)\vee A}K \rightarrow \neg(A \wedge B) \notin K)$, then $\exists K(G\mathbf{R}_{(A\wedge B)\vee A}K$ and $H = K + (A \wedge B))$.

However, since $(A \wedge B) \vee A$ is tautologically equivalent to A, we have by (**R**5):

If $G\mathbf{R}_{A\wedge B}H$ and $\forall K(G\mathbf{R}_{A}K \rightarrow \neg(A \wedge B) \notin K)$, then $\exists K(G\mathbf{R}_{A}K$ and $H = K + (A \wedge B))$.

This in turn implies (using (**R**2)):

If $G\mathbf{R}_{A\wedge B}H$ and $\forall K(G\mathbf{R}_{A}K \rightarrow \neg B \notin K)$, then $\exists K(G\mathbf{R}_{A}K$ and $H = K + (A \wedge B))$.

Using (**R**2) again, we get:

If $G\mathbf{R}_{A\wedge B}H$ and $\forall K(G\mathbf{R}_{A}K \rightarrow \neg B \notin K)$, then $\exists K(G\mathbf{R}_{A}K$ and $H = K + B)$,

that is, (**R**8).

To get (**R**7) from (**R**8), replace A by $A \vee B$ and B by A in (**R**8). We thus obtain:

If $G\mathbf{R}_{(A\vee B)\wedge A}H$ and $\forall K(G\mathbf{R}_{A\vee B}K \rightarrow \neg A \notin K)$, then $\exists K(G\mathbf{R}_{A\vee B}K$ and $H = K + A)$.

Using (**R**5), we get:

If $G\mathbf{R}_{A}H$ and $\forall K(G\mathbf{R}_{A\vee B}K \rightarrow \neg A \notin K)$, then $\exists K(G\mathbf{R}_{A\vee B}K$ and $H = K + A)$,

which is (**R**7).

Next, we derive (**R**9) from the other principles. First, substitute $A \vee B$ for A and C for B in (**R**7). We then get:

(*) If $G\mathbf{R}_{A\vee B}H$ and $\forall K(G\mathbf{R}_{A\vee B\vee C}K \rightarrow \neg(A \vee B) \notin K)$, then $\exists K(G\mathbf{R}_{A\vee B\vee C}K$ and $H = K + (A \vee B))$.

We now assume the antecedent of (**R**9), that is:

(i) $G\mathbf{R}_{A\vee B}H$ and $\neg B \in H$; and

(ii) $\forall K(G\mathbf{R}_{B\vee C}K \rightarrow \neg B \notin K)$,

in order to prove the consequent, namely: $G\mathbf{R}_{A\vee C}H$ and $\neg C \in H$.

Claim 1: $\forall K(G\mathbf{R}_{A\vee B\vee C}K \rightarrow \neg(A \vee B) \notin K)$.

In order to prove Claim 1, suppose that $G\mathbf{R}_{A\vee B\vee C}K$ and $\neg(A \vee B) \in K$. Assume also that K is inconsistent. Then, by (**R**4), $\neg(A \vee B \vee C) \in L$. This implies that also $\neg(B \vee C) \in L$. By (**R**1), there is some K′ such that $G\mathbf{R}_{B\vee C}K'$. (**R**2) yields that $B \vee C \in K'$. Hence, K′ is inconsistent and, therefore, K′ = K. Thus, $G\mathbf{R}_{B\vee C}K$. Then (ii) yields $\neg B \notin K$, contrary to the assumed inconsistency of K. Hence, K is consistent.

By the supposition and (**R**2), we have that $C \in K$. Hence, $B \vee C \in K$. K is consistent, so $\neg(B \vee C) \notin K$. Since $K = K + (B \vee C)$, we then get using (**R**1) and (**R**3), $G\mathbf{R}_{B\vee C}K$. Finally, (ii) yields $\neg B \notin K$, contrary to the supposition.

From (i), Claim (1) and (*) we conclude that there is a K such that:

$G\mathbf{R}_{A \vee B \vee C}K$ and $H = K + (A \vee B)$.

Observe that it follows from Claim (1) that K is consistent.

$\neg B \in K$ implies that $A \vee B \to \neg B \in K$. But $A \vee B \to \neg B$ is tautologically equivalent to $\neg B$, so $\neg B \in K$. This in turn implies that $K + (A \vee B) = K + A$, so $H = K + A$.

Claim 2: $G\mathbf{R}_{A \vee C}K$.

In order to prove this, we first observe that $A \vee C \in K$. This follows from the facts that $A \vee B \vee C \in K$ and $\neg B \in K$. Since K is consistent, we have that $\neg(A \vee C) \notin K$. (**R**1) and (**R**3) then yields $G\mathbf{R}_{A \vee C}K + (A \vee C)$. But $K + (A \vee C) = K$, so $G\mathbf{R}_{A \vee C}K$.

Claim 3: $\neg C \in K$.

Suppose $\neg C \notin K$. Then $\neg(B \vee C) \notin K$. Hence, $G\mathbf{R}_{B \vee C}K + (B \vee C)$. (ii) then yields $\neg B \notin K + (B \vee C)$, but this contradicts the fact that $\neg B \in K$.

It follows from Claim 2 using (**R**2) that $A \vee C \in K$. This together with Claim 3 implies that $A \in K$. Hence, $H = K + A = K$. Thus, we have proved that $G\mathbf{R}_{A \vee C}H$ and $\neg C \in H$. □

Proof of Theorem 4.5: (**R**1). If $\neg A \in L$, then $G\mathbf{R}_A\Phi$, by Definition 4.1. Suppose now that $\neg A \notin L$. Then L is an A-permitting fallback in $\mathbf{F}(\leq_G)$. By Lemma 3.13, there exists a maximal A-permitting fallback K relative to $\leq_G$. Let $H = K + A$. Then, by the definition of **R** and Definition 4.1, $G\mathbf{R}_AH$.

(**R**2) Follows from Definition 4.1 and the definition of **R**.

(**R**3) Suppose $\neg A \notin G$ and $G\mathbf{R}_AH$. All proper fallbacks are included in G (condition (**F**4)). Since G is A-permitting, G must be the *only* maximal A-permitting fallback for G. By the supposition, $H = K + A$, for some maximal A-permitting fallback for G. Hence, $H = G + A$.

(**R**4) Suppose $\neg A \notin L$ and $G\mathbf{R}_AH$. Then there is some maximal A-permitting fallback K such that $H = Cn_L(K \cup \{A\})$. But if $\bot \in H$, then $(A \to \bot) \in K$, that is, $\neg A \in K$. This is contrary to K being A-permitting. Hence, $\bot \notin H$.

(**R**5) If $\vdash_L A \leftrightarrow B$, then the maximal A-permitting fallbacks are exactly the maximal B-permitting ones. We also have that $\neg A \in L$ iff $\neg B \in L$, in this case. Hence, $G\mathbf{R}_AH$ iff $G\mathbf{R}_BH$.

(**R**6) Suppose $G\mathbf{R}_AH$ and $\neg B \notin H$. Then there is a maximal A-permitting fallback K in $\mathbf{F}(\leq_G)$ such that $H = K + A$. Suppose $\neg(A \wedge B) \in K$. Then $A \to \neg B \in K$. The deduction theorem then yields $\neg B \in H$, contrary to the supposition. Thus K is $A \wedge B$-permitting.

Assume now that $K \subset K'$ and that K′ is $A \wedge B$-permitting. Then K′ is A-permitting, contrary to the assumption that K is maximal A-permitting. Hence, K is maximal A

$\wedge$ B-permitting. But this in turn implies that $G\mathbf{R}_{A\wedge B}(K + (A \wedge B))$. But, $K + (A \wedge B) = (K + A) + B = H + B$. Hence, $G\mathbf{R}_{A\wedge B}(H + B)$.

(**R**7) Suppose that $G\mathbf{R}_A H$ and $\forall K(G\mathbf{R}_{A\vee B}K \rightarrow \neg A \notin K)$. By (**R**1), there exists a theory K_0 such that $G\mathbf{R}_{A\vee B}K_0$. The supposition, then implies that $\neg A \notin K_0$. Hence, $\neg A \notin L$. It follows that there exists a maximal A-permitting fallback K_1 in $\mathbf{F}(\leq_G)$ such that $H = K_1 + A$.

Claim: K_1 is a maximal $A \vee B$-permitting fallback.

Suppose not. Then there is a $A \vee B$-permitting fallback K_2 such that $K_1 \subset K_2$. Now, let $K = K_2 + (A \vee B)$. Then $G\mathbf{R}_{A\vee B}K$ and $\neg A \in K$, contrary to the assumption.

Let now $K = K_1 + (A \vee B)$. Then $G\mathbf{R}_{A\vee B}K$ and $H = K_1 + A = K_1 + (A \vee B) \wedge A = (K + (A \vee B)) + A = K + A$. That is, $\exists K(G\mathbf{R}_{A\vee B}K$ and $H = K + A)$.

Finally, we prove the equivalence:

($D_\leq$) $\quad A \leq_G B$ iff either $B \in L$ or $\forall H$(if $G\mathbf{R}_{\neg(A\wedge B)}H$, then $A \notin H$).

Suppose $A \leq_G B$ and that $B \notin L$. Let $G\mathbf{R}_{\neg(A\wedge B)}H$. $A \wedge B \notin L$, so H is consistent by (**R**4). It follows that there exists a maximal $\neg(A \wedge B)$-permitting fallback K in $\mathbf{F}(\leq_G)$ such that $H = K + \neg(A \wedge B)$. Assume now that $A \in H$. Then $\neg(A \wedge B) \rightarrow A \in K$. But $\neg(A \wedge B) \rightarrow A$ is tautologically equivalent to A, so $A \in K$. $A \leq_G B$ then implies by Lemma 3.6 that $B \in K$. It follows that $B \in H$. On the other hand, $\neg A \vee \neg B \in H$ and $A \in H$ imply that $\neg B \in H$, which contradicts the consistency of H. Hence, $A \notin H$. This concludes the proof of the left-to-right direction of ($D_\leq$).

For the direction from right to left, we first note that $B \in L$ implies that $A \leq_G B$, by ($\leq$2). Suppose A and B satisfy the condition $\forall H$(if $G\mathbf{R}_{\neg(A\wedge B)}H$, then $A \notin H$) and let $K \in \mathbf{F}(\leq_G)$ be such that $B \notin K$, that is, K is $\neg B$-permitting. It follows that K is also $\neg(A \wedge B)$-permitting, so, by Lemma 3.13, there is a maximal $\neg(A \wedge B)$-permitting fallback K′ in $\mathbf{F}(\leq_G)$ such that $K \subseteq K'$. Let $H = K' + \neg(A \wedge B)$. Then $G\mathbf{R}_{\neg(A\wedge B)}H$, so by the supposition, $A \notin H$. Since $K \subseteq K' \subseteq H$, it follows that $A \notin K$. We have shown that for all $K \in \mathbf{F}(\leq_G)$, if $A \in K$, then $B \in K$. Finally, Lemma 3.6 yields that $A \leq_G B$. □

Proof of Theorem 4.6: Consider any L-theory G and let $\leq$ be $\leq_G$ defined via ($D_\leq$). We proceed to prove that $\leq$ satisfies the axioms for an entrenchment ordering (cf. Definition 3.1). We first prove Dominance.

($\leq$2) Suppose that $A \vdash_L B$. We have to show that $A \leq B$, that is, either $B \in L$ or $\forall H$(if $G\mathbf{R}_{\neg(A\wedge B)}H$, then $A \notin H$). Suppose $B \notin L$ and consider any H such that $G\mathbf{R}_{\neg(A\wedge B)}H$. $\neg(A \wedge B)$ is L-consistent, since $B \notin L$. Thus, by Strong Consistency, H is L-consistent. By Success, $\neg(A \wedge B) \in H$. Suppose $A \in H$. Then $\neg B \in H$. Since $A \vdash_L B$, we also have that $B \in H$, contrary to the consistency of H. Thus, $A \notin H$.

(≤1) Suppose $A \leq B$ and $B \leq C$. We have to show that $A \leq C$. Suppose that the latter condition does not hold. Then by ($D_{\leq}$), $C \notin L$ and there is a H such that (i) $G\mathbf{R}_{\neg(A \wedge C)}H$ and $A \in H$. Since $C \notin L$, (≤2) entails that $B \notin L$. This together with $A \leq B$ implies that (ii) $\forall K$(if $G\mathbf{R}_{\neg(A \wedge B)}K$, then $A \notin K$). (i) and (ii) yield by (**R**9) that (iii) $G\mathbf{R}_{\neg(C \wedge B)}H$ and $B \in H$. (To see this substitute in (**R**9) $\neg C$, $\neg A$ and $\neg B$ for A, B and C, respectively.) Given that $C \notin L$. (iii) implies via ($D_{\leq}$) that not $B \leq C$, contrary to the assumption.

(≤3) We first prove that the following principle holds for belief revision:

(C) If $G\mathbf{R}_{A \vee B \vee C}H$ and $\neg C \in H$, then $G\mathbf{R}_{A \vee C}(H+A)$ or $G\mathbf{R}_{B \vee C}(H+B)$.

The proof of (C) mainly depends on Revision by Conjunction (**R**6). First, note that (C) would hold if H is L-inconsistent. To prove this, one uses Strong Consistency and Success. Hence, assume $G\mathbf{R}_{A \vee B \vee C}H$, $\neg C \in H$ and that H is consistent. Success and the consistency of H imply that $\neg A \notin H$ or $\neg B \notin H$. Suppose that $\neg A \notin H$. Note that $\vdash_L A \vee C \leftrightarrow [(A \vee B \vee C) \wedge (A \vee C)]$. Since $\neg A \notin H$, $\neg(A \vee C) \notin H$. Thus, applying Revision by Conjunction and Substitutivity, we get $G\mathbf{R}_{A \vee C}(H+(A \vee C))$. But $\neg C \in H$, so $H+(A \vee C) = H + A$. By a similar reasoning, we can prove that $G\mathbf{R}_{B \vee C}(H+B)$, if $\neg B \notin H$. This completes the proof of (C)

To prove (≤3), assume that $A \leq B$ and $A \leq C$ but *not* $A \leq B \wedge C$. Let us first consider the case when either $B \in L$ or $C \in L$. Then $B \wedge C$ is tautologically equivalent to either C or B. In each case, Dominance and Transitivity will imply that $A \leq B$ and $A \leq C$ entails $A \leq B \wedge C$. Therefore, suppose that neither $B \in L$ nor $C \in L$. Since we have assumed that not $A \leq B \wedge C$, ($D_{\leq}$) implies that for some H, $G\mathbf{R}_{\neg(A \wedge B \wedge C)}H$ and $A \in H$. Substituting in (C) $\neg C$ for A, $\neg B$ for B and $\neg A$ for C, we get:

If $G\mathbf{R}_{\neg C \vee \neg B \vee \neg C}H$ and $A \in H$, then $G\mathbf{R}_{\neg C \vee \neg A}(H+\neg C)$ or $G\mathbf{R}_{\neg B \vee \neg A}(H+\neg B)$,

or equivalently:

If $G\mathbf{R}_{\neg(A \wedge B \wedge C)}H$ and $A \in H$, then $G\mathbf{R}_{\neg(A \wedge C)}(H+\neg C)$ or $G\mathbf{R}_{\neg(A \wedge B)}(H+\neg B)$.

Since we have assumed the antecedent, the consequent holds, i.e., $G\mathbf{R}_{\neg(A \wedge C)}(H+\neg C)$ or $G\mathbf{R}_{\neg(A \wedge B)}(H+\neg B)$.

Since $A \in H$, $A \in H + \neg C$ and $A \in H + \neg B$. The first disjunct above therefore implies the existence of some K such that $G\mathbf{R}_{\neg(A \wedge C)}K$ and $A \in K$. This, however, contradicts the assumption that $A \leq C$ and $C \notin L$. Analogously, the second disjunct contradicts the assumption that $A \leq B$ and $B \notin L$. This concludes the proof.

(≤4) Note that Expansion has the following special case: (T) If $\bot \notin G$ and $G\mathbf{R}_T H$, then $H = G$. To prove (≤4), we first consider the case when $\bot \notin G$ and $A \notin G$. We have to show that $A \leq \bot$, i.e., that either $\bot \in L$ (which is excluded by the assumption that L is consistent). or that $\forall$H(if G$\mathbf{R}_{\neg(A\wedge\bot)}$H, then A $\notin$ H). Suppose that G$\mathbf{R}_{\neg(A\wedge\bot)}$H. Since $\neg$(A $\wedge$ $\bot$) is L-equivalent to T, (T) and Substitutivity yield that H = G. But then A $\notin$ H, since we have assumed that A $\notin$ G.

Next, we consider the case when $\bot \notin$ G and A $\leq$ $\bot$. We have to show that A $\notin$ G. Since A $\leq$ $\bot$ and $\bot \notin$ L, $\forall$H(if G$\mathbf{R}_{\neg(A\wedge\bot)}$H, then A $\notin$ H). By Seriality, there exists a theory H such that G$\mathbf{R}_{\neg(A\wedge\bot)}$H. By the same reasoning as above, H = G. Since A $\notin$ H, A $\notin$ G.

($\leq$5) Suppose that T $\leq$ A. Then by ($D_{\leq}$), either A $\in$ L, and we are done, or $\forall$H(if G$\mathbf{R}_{\neg(T\wedge A)}$H, then T $\notin$ H). But T $\in$ H, for every theory H, so the latter alternative contradicts Seriality. Hence, A $\in$ L. □

Proof of Theorem 4.7: Suppose that **R** is functional, and that, for any G, we define $\leq_G$ in terms of **R** via ($D_{\leq}$). We then define **F**($\leq_G$) from $\leq_G$ and, finally, we define **S** from **F**($\leq_G$) via ($D_{\mathbf{R}}$). We want to prove that **R** $\subseteq$ **S**, that is, for all G, H and A, if G$\mathbf{R}_A$H, then G$\mathbf{S}_A$H. It will then follow that **S** $\subseteq$ **R**. To see this, we argue as follows: Since **R** is functional, $\leq_G$ will be connected. It follows that also **S** is functional. Suppose now that G$\mathbf{S}_A$H. By Seriality, G$\mathbf{R}_A$K, for some K. **R** $\subseteq$ **S** then yields that G$\mathbf{S}_A$K. By the functionality of **S**, K = H. Hence, G$\mathbf{R}_A$H.

Suppose now that G$\mathbf{R}_A$H. We want to prove that G$\mathbf{S}_A$H. If $\neg$A $\in$ L, then, by Success, H is inconsistent. It immediately follows from the definition of **S**, that G$\mathbf{S}_A$H, in this case. Thus, we assume that $\neg$A $\notin$ L. Now, define:

$$K = \{B: \text{not}(B \leq_G \neg A)\}.$$

We want to show that:

(1) K $\in$ **F**($\leq_G$);

(2) $\neg$A $\notin$ K;

(3) for every T $\in$ **F**($\leq_G$), if K $\subset$ T, then $\neg$A $\in$ T.

(1) -(3) imply that K is a maximal A-permitting fallback in **F**($\leq_G$). Finally, we want to prove

(4) H = K + A.

Then it will follow that G$\mathbf{S}_A$H.

<u>Proof of (1)</u>: We first prove that K is closed upwards under $\leq_G$. Suppose that B $\in$ K (i.e., not(B $\leq_G$ $\neg$A)) and B $\leq_G$ C. We want to show that C $\in$ K. If C $\notin$ K, then C $\leq_G$ $\neg$A. But then, since B $\leq_G$ C and $\leq_G$ is transitive, B $\leq_G$ $\neg$A, contrary to the assumption.

Next, we prove that K is closed under conjunction. Suppose that B $\in$ K and C $\in$ K. We want to show that B $\wedge$ C $\in$ K. Suppose that this is not the case. Then B $\wedge$ C $\leq_G$ $\neg$A. However, the connectedness of $\leq_G$ yields that B $\leq_G$ C or C $\leq_G$ B. In the former case, $B \leq_G B$ (≤2) and $B \leq_G C$, so $B \leq_G B \wedge C$ (by (≤3)). In the latter case, we get $C \leq_G B \wedge C$ in a similar way. Hence, $B \leq_G B \wedge C$ or $C \leq_G B \wedge C$. But then, by transitivity, $B \leq_G \neg A$ or $C \leq_G \neg A$. This contradicts the assumption that $B \in K$ and $C \in K$.

Proof of (2): $\neg A \notin K$, since otherwise we would have that not $\neg A \leq_G \neg A$, which is impossible given the fact that $\leq_G$ satisfies Dominance.

Proof of (3): If $K \subset T$, then, for some $C \in T$, $C \notin K$. Thus, $C \leq_G \neg A$. But then, since T is closed upwards, $\neg A \in T$.

Proof of (4): We want to prove that $H = K + A$, i.e.,

(i) $B \in H$ iff $A \rightarrow B \in K$.

To prove (i) it is sufficient to prove:

(ii) $B \in H$ iff not$(A \rightarrow B \leq_G \neg A)$.

Now, by ($D_{\leq}$) and the fact that $\neg A \notin L$, and the fact that **R** is functional,

(iii) not$(A \rightarrow B \leq_G \neg A)$ iff $A \rightarrow B \in T$,

where T is the (unique!) **R**-revision of G with $\neg((A \rightarrow B) \wedge \neg A)$. Now, $\neg((A \rightarrow B) \wedge \neg A)$ is L-equivalent to A. Hence, by Substitutivity, T is the **R**-revision of G with A. But then, $T = H$. Thus, to prove (ii), it is sufficient to prove:

(iv) $B \in H$ iff $A \rightarrow B \in H$.

But this is trivial in view of the fact that $A \in H$ (by Success). □

Proof of Lemma 5.1: We first consider the limiting case, in which $\neg A \in L$. In this case, (a) $G\mathbf{R}_A H$ iff $\perp \in H$ (by Success and Seriality) and (b) $G\mathbf{R}^{-}_{\neg A} K$ iff $G = K$, by (C3).

To prove *Levi* from right to left, assume that $G\mathbf{R}_A H$. Then, by (a), $\perp \in H$. Therefore, since $\neg A \in L$, $H = G + A$. Also, by (b), $G\mathbf{R}^{-}_{\neg A} G$. Hence, the r. h. s. of *Levi* is satisfied for $K = G$.

For the other direction, suppose that for some K, $G\mathbf{R}^{-}_{\neg A} K$ and $H = K + A$. Since $\neg A \in L$, $\perp \in H$. But then (a) implies that $G\mathbf{R}_A H$.

We now consider the principal case, in which $\neg A \notin L$. We first prove *Levi* from left to right. Suppose $G\mathbf{R}_A H$. Then, by definition, there is some maximal A-permitting fallback K such that $H = K + A$. By (C2), there is some theory T such that $G\mathbf{R}^{-}_{\neg A} T$ and $K \subseteq T \subseteq G \cap H$. We want to show that $H = T + A$. Since $K \subseteq T$, the monotonicity of expansion yields $K + A \subseteq T + A$, that is, $H \subseteq T + A$. Since, $T \subseteq H$ and $H + A = H$, monotonicity yields $T + A \subseteq H$.

For the other direction of *Levi*, suppose that for some K, $G\mathbf{R}^{-}_{\neg A} K$ and $H = K + A$. By (C1), there is a maximal A-permitting fallback T such that $T \subseteq K \subseteq G \cap (T + A)$.

By the definition of revision, $G\mathbf{R}_A(T + A)$. Thus, $G\mathbf{R}_A H$ will follow if we can prove that $K + A = T + A$. Since $T \subseteq K$, $T + A \subseteq K + A$. Since $K \subseteq T + A$, $K + A \subseteq (K + A) + A = K + A$. □

NOTES

* The inspiration for this paper comes from our discussions with the participants of the Workshop on the Logic of Theory Change — Konstanz, October 13 - 15, 1989. We want to thank them all. Special thanks go to the organizers of the conference André Fuhrmann and Michael Morreau. We are also grateful to Sven Ove Hansson and Hans Rott for their very helpful comments on an earlier version.

[1] Department of Philosophy, University of Uppsala, Villavägen 5, S-752 36 Uppsala, Sweden.

[2] Actually, Grove works with an ordering of epistemic *plausibility*. But as Gärdenfors (1988, sect. 4.8) points out, the notions of plausibility and entrenchment are interdefinable. Thus, a proposition A is at least as plausible as a proposition B given the agent's beliefs if and only if non-B is at least as entrenched as non-A in the agent's belief set. The notion of epistemic entrenchment is primarily defined for the propositions that belong to the agent's belief set: one adopts the convention that propositions that are not believed by the agent are minimally entrenched. On the other hand, the notion of plausibility primarily applies to the propositions that are incompatible with the agent's beliefs (the propositions that are compatible with what he believes are all taken to be equally and maximally plausible). Thus, this is a notion of *conditional* plausibility. A is at least as plausible as B in this sense iff the following holds: on the condition that I would have to revise my beliefs with $A \vee B$, I should change them in such a way as to allow for A.

[3] Cf. our informal characterization of epistemic entrenchment above.

[4] For the notion of a closure system and other basic notions from universal algebra and lattice theory that appear in this section, see for example Grätzer (1979), Chap 1 and Rasiowa & Sikorski (1970), Chap 1.

[5] Rott assumes connectedness for entrenchment. Therefore, he works with nested families of fallbacks, or, as he calls them, "EE-cuts".

[6] Cf. Rott (1989), for the same observation.

REFERENCES

Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985) 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic* **50**, 510-530.

Gärdenfors, P. (1988) *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press.

Grove, A. (1988) 'Two modellings for theory change', *Journal of Philosophical Logic* **17**, 157-170.

Grätzer, G. (1979) *Universal Algebra* (Second Edition), Springer Verlag.

Levi, I. (1980) *The Enterprise of Knowledge,* the MIT Press.

Lindström, S. and Rabinowicz, W. (1989) 'On Probabilistic Representation of Non-Probabilistic Belief Revision', *Journal of Philosophical Logic* **18**, 69-101.

Lindström, S. and Rabinowicz, W. (1990) 'Belief Revision, Epistemic Conditionals and the Ramsey Test', *Uppsala Philosophy Reports* no 1, 1990.

Niederée, R. (1990) 'Multiple Contraction. A Further Case Against Gärdenfors' Principle of Recovery', this volume.

Rasiowa, H., and Sikorski, R. (1970) *The Mathematics of Metamathematics* (Third edition), PWN —Polish Scientific Publishers, Warsaw.

Rott, H. (1989) 'Two Methods of Contractions and Revisions of Knowledge Systems'. In Michael Morreau (ed.), *Proceedings of the Tübingen Workshop on Semantic Networks and Non-Monotonic Reasoning,* Vol. 1, SNS-Bericht 89-48, Tübingen, pp. 28-47.