"Just Because": Taking Belief Bases Seriously

Hans Rott

University of Amsterdam, ILLC/Department of Philosophy
Nieuwe Doelenstraat 15, 1012 CP Amsterdam, The Netherlands
rott@philo.uva.nl

Abstract. We formalize several ways of accounting, in the context of logically closed theories, for foundationalist intuitions that underlie change operations applying to belief bases. A positive and a negative concept of entrenchment is defined on the basis of the structure of a given, possibly prioritized belief base. Only the latter, more fine-grained concept proves to be appropriate for a successful attempt at approximating base changes on the theory level. We investigate the question as to which degree we can comply with the fundamental intuition expressed by the various Filtering Conditions that say that all (and only) beliefs that are believed "just because" a retracted belief was believed should be withdrawn.

1 Introduction

The problem dealt with in the present paper is best illustrated by an example.

Example 1. Consider a theory $K = Cn(\phi, \neg\psi)$ which we want to contract with respect to ϕ. Assume that K is generated by the belief base $H = \{\phi, \neg\psi\}$ and that ϕ enjoys epistemic priority over $\neg\psi$. It seems intuitively clear that $\neg\psi$ should be in the contracted belief set and in fact that this set should be identified with $Cn(\neg\psi)$. According to a model proposed by Gärdenfors and Makinson (1988), this can be so if and only if the disjunction $\phi \vee \neg\psi$, which is not itself included in the base H, is more entrenched than ϕ alone. Since the disjunction is backed by both elements of the belief base this appears to be a sensible idea (see Figure 1).

From the informal argument just given, it might seem that the entrenchment of a derived belief varies with the number of basic beliefs ("premises") supporting it. However, this idea does not match well with current theories about comparative epistemic entrenchment.[1] Additional problems for *combinatorial modellings* of entrenchment might arise when the elements of a belief base are not logically independent. What we need in this situation is some systematic way of extending the orderings of the elements in the belief base into orderings of the whole

[1] For example, the Gärdenfors and Makinson's (1988) conjunctiveness condition, according to which the entrenchment of a conjunction cannot be less than the entrenchment of both of its conjuncts, fails.

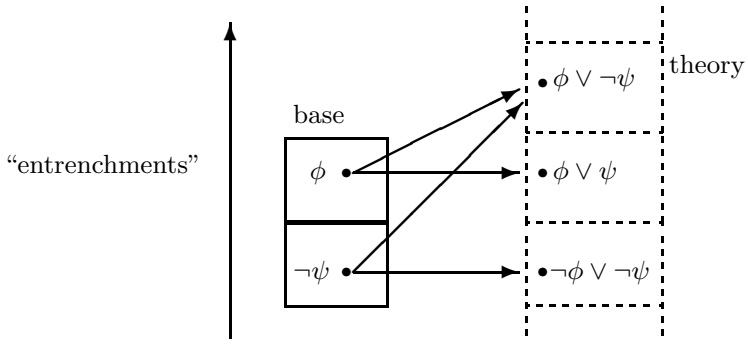

Fig. 1. How to generate entrenchments from base priorities?

generated theory. What is the right way of arguing for the entrenchments of the elements of a theory that has been generated by some prioritized knowledge base?

This paper is about the relation between belief/knowledge *base* revision and *theory* revision. The former tends to be more interesting for computer scientists, the latter for logicians and philosophers. Our question is how the former can be "rationally reconstructed" with tools from the latter. More specifically, our program is as follows. We start with some general remarks about the interpretation of a belief base containing basic bits of information ordered by some priority relation. The underlying philosophy of base revision has a foundationalist flavour.[2] We look at one necessary condition for respecting the foundationalist intuitions at "the knowledge level" of theory change, viz., the so-called Filtering and Simple Filtering conditions. These conditions that tell us not to keep beliefs that are essentially dependent on jettisoned beliefs, come in two versions, according to whether we understand the phrase "just because" as referring only to basic beliefs, or to derived beliefs as well.

Next we construct and contrast two different conceptions of epistemic entrenchment which are relevant in the context of prioritized base revisions. There is a "positive" relation, $<^+$, which is a numerically representable relation of epistemic entrenchment in the standard sense of Gärdenfors and Makinson (1988). The construction of $<^+$ has been quite common in earlier literature. Then we introduce a "negative" relation, $<^-$, which is closer to the intuitions bound up with the term "entrenchment." However, $<^-$ is only an entrenchment relation in a liberalized sense (see Rott 1992b). Due to incomparabilities, it is not numerically representable, and it almost unknown in the existing literature. This negative entrenchment relation turns out to be a refinement of the positive relation.

[2] Unfortunately, the epistemological distinction between foundationalism and coherentism cannot be explained here. For these notions and suggestions of how to explicate them in terms of models of belief revision, see Gärdenfors (1990), Doyle (1992), Rott (1996), del Val (1997) and Olsson (1997).

Then we recall two ways of using entrenchment relations in theory change, the standard Gärdenfors-Makinson recipe and Rott's more "brutal" (Williams 1994) recipe. Neither of these can be sensibly used in conjunction with the positive relation $<^+$. The negative relation is more suitable. If we use the Gärdenfors-Makinson recipe for entrenchment-based contraction with $<^-$, then we satisfy the condition of Simple Filtering referring to basic beliefs, but not the stronger one referring to all beliefs. If we use the recipe for brute entrenchment-based contraction with $<^-$, then we even satisfy the condition of Simple Filtering referring to all beliefs. We recall from Rott (1992c) that the two recipes can serve as upper and lower bounds featuring in Lindström-Rabinowicz's (1991) interpolation thesis, and we give an argument that appears to show that no better approximation of prioritized base contraction on the knowledge level is possible, at least if we confine ourselves to reconstructions in terms of epistemic entrenchment.

In the last section we show that a quantitative contraction mechanism introduced by Mary-Anne Williams suffers from the so-called "drowning effect". The basic reason for this is that it essentially uses an ordinalized version of the relation $<^+$ in conjunction with the Gärdenfors-Makinson contraction recipe – a combination we shall already find inadequate in the case of the qualitative notion of entrenchment.

Formalities. We are working in a propositional language L with variables p, q, r etc. We use ϕ, ψ, χ etc. as metavariables for sentences from L, and G, H, K etc. as metavariables for subsets of L. Curly letters like $\mathcal{H}$ indicate more complex structures. The background logic Cn (consequence operation) alias $\vdash$ (consequence relation) governing L is supposed to be Tarskian (reflexive, monotonic, transitive and compact). For the sake of simplicity, we are going to think of it as classical propositional logic.

2 Belief bases with priorities

Hansson (1994) has argued that we should take belief bases seriously. His main point was that only on belief bases can we perform local changes; logically closed theories are always so large and tightly inferentially interwoven that local changes are impossible. In this paper, we shall give more, alternative interpretations of why it is important to apply change operations to belief bases rather than to fully-fledged theories.

Elementship in a belief base carries important epistemological implications.

Maxim B The elements in a belief base are basic (fundamental, explicit) beliefs. They comprise beliefs, and only beliefs, which have some kind of independent standing, i.e., which are not derived from other beliefs.

The elements in a belief base have a distinguished status. Each of them is a singular piece of information, irrespective of the syntactical complexity of the sentences expressing them. The difference between one item $p \wedge q$ and two items

p and q in a knowledge base is not just "notational bondage" that should be straightened out by some process of "articulation". These are Belnap's (1979, p. 23) words. I think that Rescher in his reply to Belnap is entirely correct in pointing out that if we take belief bases seriously (e.g., for the purpose of hypothetical reasoning), then there is indeed a crucial difference between "juxtaposing commas" and "conjoining ampersands" (1979, p. 31). Belnap, however, keeps on denying the difference even in the amended version of his critique published in Anderson, Belnap and Dunn (1992, pp. 541–553). Other work that warns us not to thoughtlessly form conjunctions can be found among paraconsistent logicians (Jaskowski 1948/1969, Rescher and Brandom 1979, Urchs 1995) and researchers interested in the interface of logic and probability (Kyburg 1970, 1997).

Evidently, a belief base is not just an axiomatization of a theory. For instance, we must not make belief bases non-redundant, something that we certainly want to do for axiomatic systems. According to Maxim B, it is not legitimate to subject the belief base to *any* manipulations. While eliminating redundancies, or closing under conjunctions or disjunctions or under some non-classical logic may have desirable effects from a formal point of view, it runs counter to Maxim B. It is the structure of the original belief base that provides us with information about how to change our beliefs, and that structure should not be changed without compelling reasons.

Intuitively, derived beliefs are believed only *because* basic beliefs are believed.[3] In real applications, items of basic or explicit belief are often associated with some degree of certainty. Formally, therefore, we model a *prioritized belief base* as a sequence of sets of sentences

$$\mathcal{H} \;=\; \langle H_1, \ldots, H_n \rangle$$

For $i < j$, the elements in H_j are more "important" or "reliable", they have more "weight" than the elements in H_i. Given a prioritized belief base $\langle H_1, \ldots, H_n \rangle$, we denote the set $\bigcup_{i=k}^{n} H_i$ by $H_{\geq k}$. The flat or simple base

$$H \;:=\; |\mathcal{H}| \;:=\; H_{\geq 1}$$

generates a theory (at the "knowledge level")

$$K \;:=\; Cn(H)$$

The theory K may be thought of as comprising both basic and derived (explicit and implicit) beliefs.

This representation of a belief base in terms of a sequence of sets is somewhat more general than conceiving of a prioritized base as a set H equipped with a

[3] This "because" is of a rational, not of a causal character. Bernard Williams's (1970, pp. 100–102) argument that a rational "because" between beliefs is in general also a causal "because" does not strike me as convincing.

weak ordering $\prec$,[4] since the H_i's in $\mathcal{H}$ need not be disjoint. We may have one occurrence of ϕ in H_i and one in H_j, with i different from j, something that would not be possible in the alternative formalization using $\prec$. For the purposes of this paper, however, one can drop all non-maximal occurrences of sentences without any change of the results. So we suppose in what follows that the H_i's in $\mathcal{H}$ are actually disjoint, and we can equivalently speak of the set H equipped with a priority ordering $\prec$, with the understanding that for any two elements ϕ and ψ in H, we have $\phi \prec \psi$ if and only if $\phi \in H_i$ and $\psi \in H_j$ for i and j with $i < j$. Clearly, the flat belief base $\mathcal{H} = \langle\{H\}\rangle$ is associated with the empty ordering $\prec$.

It is worth emphasizing that each prioritized belief base conveys two kinds of information. First, we have syntactical information that distinguishes, e.g., the joint belief $\phi \wedge \psi$ from two simultaneous but separate beliefs ϕ and ψ. In the former case, ϕ and ψ will stand and fall together while in the latter case they are going to be treated as independent. Second, of course, the prioritized belief base acknowledges varying degrees of belief, with the idea that beliefs in higher layers should in case of conflict be given precedence over beliefs in lower layers of the belief base $\mathcal{H}$.

Sometimes one may wish to abstract from the two kinds of structure. The priority structure is lost in the move from $\mathcal{H}$ to H, the syntactic structure is lost in the move from H to K.[5] But although K *itself* does not record its origin in $\mathcal{H}$, one may suppose that potential *changes* of K are dependent on its "provenance" or "history". Our central question now is: *How* can the structure of the prioritized belief base $\mathcal{H}$ guide the changes made to the generated theory K which in itself does no longer itself reflect the relevant structure? We propose to use relations of epistemic entrenchment as a tool for the construction and reconstruction of theory revisions.

3 Filtering and simple filtering

If we want to take belief bases seriously, we should somehow respect the argumentative dependencies of our beliefs when performing belief changes. One proposal to capture this idea is the *Filtering* condition of Fuhrmann (1991):[6]

(F) When a belief ψ is withdrawn in order to form the contraction of one's beliefs by ϕ, then the contraction by ϕ should not contain any sentences that were believed "just because" ψ was believed.

[4] An asymmetric *weak ordering* is a relation $\prec$ that is (a) modular in the sense that $x \prec y$ entails that either $x \prec z$ ore $z \prec y$, for all x, y and z (transitivity follows from this), but that is (b) not antisymmetric, that is, it does not follow from the fact that x and y stand in the same $\prec$-relation with every z that x and y are identical.

[5] Or in the move from H to $\bigwedge H$; compare Nebel (1989) and del Val (1997).

[6] Also compare Martins and Shapiro's (1988) "disbelief propagation", and Rao and Foo's (1989) "foundational equation", according to which foundationalism is coherentism plus disbelief propagation.

This is quite a complex requirement. In the present paper, we are going to study only a much simpler version of the Filtering condition, which we call the condition of *Simple Filtering*.

(SF) The contraction of one's beliefs by ϕ should not contain any sentences that were believed "just because" ϕ was believed.

Notice that the more radical a contraction operation is in removing beliefs, the more likely it is to satisfy the Filtering requirement (in either of its versions). Against the ideology that is often associated with the belief revision paradigm initiated by Alchourrón, Gärdenfors and Makinson, (Simple) Filtering favours non-conservative or non-minimal incisions into one's belief state. It can at most be a necessary, but not a sufficient condition for the rationality of belief change.

The statements of (F) and (SF) have deliberately been left vague. Evidently, those sentences that are believed just because some other sentence is believed are derived (implicit) beliefs. But the conditions are ambiguous with regard to the status of the sentences that are to be retracted in the first place. Let us distinguish two versions of (SF).

(BSF) The contraction of one's beliefs by ϕ should not contain any sentences that were believed "just because" ϕ was a basic belief.

(DSF) The contraction of one's beliefs by ϕ should not contain any sentences that were believed "just because" ϕ was a (basic or derived) belief.

Intuitively, (BSF) seems too restricted a condition and the more demanding condition (DSF) is therefore to be preferred.

4 "Just because"

Now we need to know what the phrase "just because" means. Corresponding to the two different versions of Simple Filtering, we distinguish two variant readings of "just because". The first one applies to the elimination of *basic* beliefs ϕ:

(BJB) A sentence ψ is in $K = Cn(H)$ *just because* ϕ is in H, if ψ is in K but not in $Cn(H - \{\phi\})$.[7]

According to this definition, ψ is in K just because ϕ is in H, only if ϕ is in fact in H and H is not logically closed. If H is logically closed and contains ϕ, then both $\phi \vee \chi$ and $\phi \vee \neg\chi$ are in $H - \{\phi\}$, so set-theoretically removing ϕ from H does not affect the consequences of H at all.

An equivalent way of expressing the above definition is to say that ψ is in K "just because" ϕ is in H if and only if every way of deriving ψ draws on ϕ. It is presupposed here that ϕ is not only in the theory K but actually in the belief

[7] This definition is the same as Fuhrmann's (1991, p. 185) definition of χ's being "H-dependent on ψ."

base H. Removing this presupposition, we can say more generally that ψ is in K "just because" ϕ is in K if and only if every way of deriving ψ also derives ϕ. There is no way of deriving ψ without at the same time deriving ϕ. Our second reading of "just because" applies to the elimination of *derived* beliefs ϕ as well:

(DJB) A sentence ψ is in $K = Cn(H)$ *just because* ϕ is in K, if ϕ is not in $Cn(\emptyset)$ and ψ is in K but not in $Cn(G)$, for all $G \subseteq H$ such that ϕ is not in $Cn(G)$.

According to this definition, ψ is in K just because ϕ is in K if the base offers no way of keeping ψ while getting rid of ϕ. Of course, this can only happen if ϕ is in fact in K, but ϕ need not be in the base H. In the following it is understood that (BSF) uses the reading (BJB) and that (DSF) uses (DJB).[8]

Observation 1. If ψ is in the theory $K = Cn(H)$ just because ϕ is in the belief base H, then ψ is in K just because ϕ is in K. The converse is not valid, even when ϕ is in H. As a consequence, (DSF) implies (BSF), but not conversely.

Proof. Let ψ be in the theory K just because ϕ is in the base H. By definition, this means that ψ is in $K = Cn(H)$ and that ψ is not in $Cn(H - \{\phi\})$. Hence ϕ is in H, and ϕ is not in $Cn(\emptyset)$. By the reflexivity of Cn, ϕ is in $K = Cn(H)$. It remains to show that

$$\text{for every } G \subseteq H \text{ such that } G \nvdash \phi \text{ it does not hold that } G \nvdash \psi$$

But clearly, if $G \subseteq H$ with $G \nvdash \phi$, we have $G \subseteq H - \{\phi\}$. Since ψ is not in $Cn(H - \{\phi\})$, we get $G \nvdash \psi$, by the monotonicity of Cn. Therefore ψ be in the K just because ϕ is in K.

Counterexample to the converse direction. If we have a redundant base like $H = \{p, p \supset q, q\}$, then $q \vee r$ is *not* in K just because q is *in* H, but $q \vee r$ *is* in K just because q is *in* K.

Since (BJB) implies (DJB) and the corresponding concepts figure in the antecedents of (BSF) and (DSF), it follows immediately that (DSF) implies (BSF). That the converse is invalid will be seen from Observation 5. QED

The above counterexample shows that the two readings of "just because" make a difference. The result depends on whether we refer to q as an explicit item in the belief base or as a derived belief in the generated theory. Intuitively (DJB), the definition referring to explicit as well as implicit beliefs, seems to be more adequate than (BJB).

[8] Now formally, (DSF) becomes

$$\text{If } [\text{ for every } G \subseteq H, \text{ if } G \nvdash \phi \text{ then } G \nvdash \psi\,], \text{ then } \psi \notin K \dot{-} \phi$$

One might think that this condition is trivially satisfied, since we certainly require that $K \dot{-} \phi \nvdash \phi$. However, notice that the antecedent refers to subsets *of* H while $K \dot{-} \phi$ is only a subset *of* K.

An important problem left over from our discussion of Example 1 is this: How can the syntax and the prioritization of the base H generate a well-balanced entrenchment relation that can be construed as guiding revisions of the the full theory K?

5 Comparing the quality of subsets of prioritized belief bases

We assume as given a prioritized belief base $\mathcal{H} = \langle H_1, \ldots, H_n \rangle$ and make use of the corresponding priority relation $\prec$ over $H = |\mathcal{H}|$. We introduce two different ways of comparing the quality of the subsets of H.

5.1 The positive way

Subsets G and G' of the belief base H may be taken to be *premise sets* for deriving further beliefs. In this role, they can be compared by looking essentially at their *weakest elements.*

Definition 1. $G \lhd^+ G'$ iff G is non-empty and for every χ in G' there is a ξ in G such that $\xi \prec \chi$.

Smaller sets tend to be better (i.e., greater) with respect to $\lhd^+$ than bigger sets. For some properties of $\lhd^+$ and how they depend on the properties of $\prec$, see Rott (1992a, pp. 126–128).

We can introduce numbers representing the quality of subsets G of H, considered as premise sets. Let, for subsets G of H, $\text{qual}^+(G) = \min\{i : G \cap H_i \neq \emptyset\}$, with the convention that $\min(\emptyset) = n + 1$. Then we can equivalently say that $G \lhd^+ G'$ iff $\text{qual}^+(G) < \text{qual}^+(G')$.

5.2 The negative way

Subsets G and G' of the belief base H may alternatively be taken to be *remainder sets* consisting of the pieces of information one may retain when forced to give up some belief. In this role, they are compared by looking essentially at their *strongest non-elements*, or more exactly, by the *strongest* elements of H which are *missing* from one of them but not from the other.

Definition 2. $G \lhd^- G'$ iff $G \neq G'$ and for every χ in $G - G'$ there is a ξ in $G' - G$ such that $\chi \prec \xi$.

We can equivalently say that $G \lhd^- G'$ iff for $k = \max\{i : (G \Delta G') \cap H_i \neq \emptyset\}$ it holds that $(G \Delta G') \cap H_k$ is entirely contained in G'. (Δ is the symmetric difference operator.)

Or, again equivalently, we can say that $G \lhd^- G'$ iff there is some i such that $G \cap H_i \subset G' \cap H_i$ and for all $j > i$, we have $G \cap H_j \subseteq G' \cap H_j$.

Bigger sets tend to be better (greater) with respect to $\lhd^-$ than smaller sets. For some properties of $\lhd^-$ and how they depend on the properties of $\prec$, see Rott (1992c, pp. 30–33).

Again, we can introduce numbers representing the quality of subsets G of H, this time considered as remainder sets. Let, for subsets G of H, $\text{qual}^-(G) = n + 1 - \max\{i : H_i - G \neq \emptyset\}$, with the convention that $\max(\emptyset) = 0$. Then we can equivalently say that $G \lhd^- G'$ iff $\text{qual}^-(G) < \text{qual}^-(G')$.

Keep in mind that the positive and the negative qualities of sets of explicit beliefs are very different concepts. Whereas a singleton set taken from H_n is very good as a premise set, it is in general very bad as a remainder set; the reverse is true of the flattened base H itself.

6 Defining entrenchments in theories generated by prioritized belief bases

We are now going to define the entrenchment of sentences which are derived from a given prioritized belief base $\mathcal{H} = \langle H_1, \ldots, H_n \rangle$. We do this in two ways, by using the two relations introduced on the last section.

6.1 The positive way

We can compare the entrenchment of two elements in a theory generated by a prioritized belief base in a positive way, by comparing all potential ways of *deriving* the beliefs in question, that is, by comparing all potential *premise sets implying* the beliefs in question.[9]

Definition 3. ϕ is more entrenched$^+$ than ψ, in symbols $\psi <^+ \phi$, iff ϕ is implied by H and for every premise set $G \subseteq H$ implying ψ there is a premise set $G' \subseteq H$ implying ϕ such that $G \lhd^+ G'$.

This positive way of thinking of entrenchments has been quite popular in the literature. With minor variations, it has been suggested by Rescher (1964, pp. 49–50; 1976, pp. 18–19), Dubois and Prade (1991, Section 4; 1992, Section 3), Rott (1991b, Section 3; 1992a, Section 4, here without a connectedness presumption for $\prec$) and Williams (1994, Section 4).

The positive generation mechanism should be viewed as a means of constructing an entrenchment relation from a finite representation of it, viz. from a prioritized belief base (or more exactly, from an "E-base" satisfying a certain entailment condition, see Rott 1991b, pp. 142–146). As we shall see, the positive entrenchment relation is not suitable for application in entrenchment-based theory contraction in the style of Gärdenfors and Makinson (1988) or Rott (1991a).

[9] In Section 5.1 of Rott (1992c), I rashly dismissed positive relations of entrenchment, since there (a) I only considered them in conjunction with the relation $\lhd^-$ which is not suitable in this context, (b) I only looked at flat belief bases, and (c) I did not see how a positive relation might be used for entrenchment-based contractions.

For this purpose, a negative relation of entrenchment turns out to be much more appropriate.

In a theory generated by a flat belief base, all elements except the logical truths are equally entrenched in the sense of $<^+$. We take down the following

Observation 2. $\psi <^+ \phi$ iff there is a k such that $H_{\geq k} \vdash A$ but not $H_{\geq k} \vdash \psi$.

In order to prove this observation, we first introduce a useful concept. Given some prioritized belief base $\mathcal{H}$, the entrenchment of a sentence ϕ (with respect to $\mathcal{H}$), in symbols $\text{ent}^+(\phi)$, is the maximal number k such that $\phi \in Cn(H_{\geq k})$. Equivalently, $\text{ent}^+(\phi) = \max\{\text{qual}^+(G) : G \subseteq H \text{ and } G \vdash \phi\}$. If ϕ is not in $Cn(H)$, we put $\text{ent}^+(\phi) = 0$.

Proof of the Observation. Let $\psi <^+ \phi$. Take $k = \text{ent}^+(\phi)$. By definition, $H_{\geq k} \vdash \phi$. Suppose for reductio that also $H_{\geq k} \vdash \psi$. Since $\psi <^+ \phi$, there must be a set $G' \subseteq H$ implying ϕ such that $H_{\geq k} \lhd^+ G'$. But any set which is $\lhd^+$-better than $H_{\geq k}$ must be a subset of $H_{\geq k+1}$, so $H_{\geq k+1} \vdash \phi$, contradicting the choice of k.

For the converse, let there be an i such that $H_{\geq i} \vdash \phi$ but not $H_{\geq i} \vdash \psi$. Then all premise sets $G \subseteq H$ implying ψ must contain elements from $H - H_{\geq i}$. Thus for all premise sets $G \subseteq H$ implying ψ it holds that $G \lhd^+ H_{\geq i}$, so $\psi <^+ \phi$. QED

It is now easy to realize that the positive entrenchment relation $<^+$ is *numerically representable* by the ordinal entrenchment function ent^+ in the sense that $\phi <^+ \psi$ if and only if $\text{ent}^+(\phi) < \text{ent}^+(\psi)$.

6.2 The negative way

We can compare the entrenchment of two elements in a theory generated by a prioritized belief base in a negative way, by comparing all potential ways of *discarding* the beliefs in question, that is, by comparing all potential *remainder sets not implying* the beliefs in question.

Definition 4. ϕ is more entrenched$^-$ than ψ, in symbols $\psi <^- \phi$, iff ψ is not a logical truth and for every remainder set $G \subseteq H$ that does not imply ϕ there is a remainder set $G' \subseteq H$ that does not imply ψ and $G \lhd^- G'$.

This relation is supposed to capture a "minimal change interpretation" of the notion of entrenchment which coincides, in the context of prioritized base changes, with a "competitive interpretation". It is not a relation of epistemic entrenchment in the standard sense, but only in a generalized sense. (For all this, see Rott 1992c, pp. 36–45). There is no ordinal function ent^- representing the relation $<^-$, since the relation $<^-$ will in general present us with incomparabilities of beliefs in terms of entrenchment.

Example 2. Let $H = \{p, p \vee q, p \supset q\}$. Without any knowledge about the prioritization of H, we can say that the set of all $\lhd^-$-best subsets of H not

implying p is $\{\{p \vee q, \neg p \vee q\}\}$, of those not implying q is $\{\{p, p \vee q\}, \{\neg p \vee q\}\}$, and of those not implying $p \vee q$ is $\{\{\neg p \vee q\}\}$. This gives us $p <^- p \vee q$, but neither $p <^- q$ nor $q <^- p \vee q$. That means that q is not comparable in terms of $<^-$ with either p or $p \vee q$. – Notice by the way that prioritization and the negative concept of entrenchment are two altogether different kinds of thing. For if the pairs related by the prioritization $\prec$ of H were exactly $p \vee q \prec p$ and $p \vee q \prec p \supset q$, then we would have $p <^- p \vee q$ although $p \vee q \prec p$ (and of course neither $p \vee q <^- p$ nor $p \prec p \vee q$). The reflective equilibrium reached by an assessment of entrenchments may strictly reverse the prima facie evaluation of basic beliefs in terms of their priorities.

Although the "positive" quality of a set of explicit beliefs (as a premise set) is very different from its "negative" quality (as a remainder set) the positive and negative concepts of entrenchment are closely connected.

Observation 3. The negative relation $<^-$ of epistemic entrenchment refines the positive one, $<^+$.

Proof. Let $\phi <^+ \psi$. We have to show that $\phi <^- \psi$, that is, we have to show that for all remainder sets $G \subseteq H$ which do not imply ψ there is a remainder set $G' \subseteq H$ with $G \lhd^- G'$ which does not imply ϕ. Take an arbitrary remainder set $G \subseteq H$ that does not imply ψ. Using Observation 2, we conclude from $\phi <^+ \psi$ that there is a k such that $H_{\geq k}$ implies ψ but not ϕ. Now define $G' = H_{\geq k}$. It remains to show that $G \lhd^- H_{\geq k}$. That is, we must show that for each χ in $G - H_{\geq k}$ there is a ξ in $H_{\geq k} - G$ such that $\chi \prec \xi$. By the definition of $H_{\geq k}$, we know that $\chi \prec \xi$ for all such χ and all ξ in $H_{\geq k}$. Thus it remains to show that $H_{\geq k} - G$ is non-empty. But this follows from the fact that $H_{\geq k}$ implies ψ and G does not imply ψ (and Cn is monotonic).

The converse is not valid. We give a simple example where $\phi <^- \psi$ but not $\phi <^+ \psi$. Consider the flat belief base $H = \{p, q\}$. Since $\emptyset \lhd^- \{q\}$, we get that $p <^- p \vee q$, but clearly not $p <^+ p \vee q$. QED

The positive relation $<^+$ generates "layers" of beliefs through the definition of equivalences, putting $\phi \sim^+ \psi$ iff neither $\phi <^+ \psi$ nor $\psi <^+ \phi$. The negative relation $<^-$ introduces new distinctions within these layers, and in general it generates *incomparabilities* within these layers. But since it is a refinement of $<^+$, the coarser layer structure of the latter remains intact.

Clearly, the above ways of constructing entrenchment relations are not the only conceivable ones. For example, one might try the use of difference sets in the positive approach to entrenchment, or to avoid the use of difference sets in the negative approach. We believe, however, that we have hit on two particularly important constructions.

7 How to use prioritized belief bases in belief change

We describe a natural method of theory change guided through the structure of a prioritized belief base. The principal idea of the method of *prioritized belief*

base contraction is the following:

$$\psi \in K\dot{-}\phi \quad \text{iff} \quad \psi \in Cn(G), \text{ for every } \lhd^{-}\text{-best } G \subseteq H \text{ such that } G \not\vdash \phi$$

According to this recipe, a sentence ψ remains in the contraction of K by ϕ just in case it follows from every best remainder set of the belief base that does not imply ϕ. This idea has been studied, among others, by Rescher, Fagin/Ullman/Vardi, Brewka, Nebel, Rott, del Val and several researchers from Toulouse. More methods of exploiting the structure for prioritized belief bases in processes of belief revision or nonmonotonic reasoning are surveyed in Cayrol and Lagasquie-Schiex (1995).

8 How to use entrenchment relations in belief change

In this section, we assume as given some relation $<$ of epistemic entrenchment over a theory K. The official Gärdenfors-Makinson (1988) construction recipe for *theory* contraction based on epistemic entrenchment is this.

(Def$\dot{-}$) $$\psi \in K\dot{-}\phi \quad \text{iff} \quad \psi \in K \text{ and } \phi < \phi \vee \psi$$

For the motivation of this definition and the use of the disjunction (which looks counterintuitive at first sight), see Gärdenfors and Makinson (1988, pp. 89–90) and Rott (1992c, Section 7).

Is this recipe applicable in the context of belief *base* contraction? In general, a disjunction $\phi \vee \psi$ is not included in the belief base H, even if both ϕ and ψ are in H. So one idea is that one should "blow up" belief bases by closing it under disjunctions.[10] However, with such a move one clearly departs from the philosophy underlying the belief base approach as formulated in Maxim B. It seems to us that there is no good intuitive justification for stipulating that belief bases should be disjunctively closed. If we begin to process the original collection of data at all, then why not take into consideration *everything* that can be derived from it, that is, the *whole* theory generated by the data base? But this would mean dismissing the idea of belief base change in favour of theory change.

There is another proposal of how to use epistemic entrenchment that is not afflicted with the same problems. Rott (1991a) ventilates the following method which is further studied in Fermé and Rodriguez (1998) and Pagnucco and Rott (1998).

(Def$\ddot{-}$) $$\psi \in K\ddot{-}\phi \quad \text{iff} \quad \psi \in K \text{ and } \phi < \psi$$

This method is very rude, however, since it eliminates many weakly entrenched beliefs in the contraction with respect to ϕ, even if they stand in no substantial connection with ϕ. Due to the properties of entrenchment relations,

[10] This idea seems to motivate Hansson (1993) and Williams (1994; 1995).

$K \ddot{-} \phi$ is always a subset of $K \dot{-} \phi$. For this reason we may call $K \ddot{-} \phi$ the *severe* and $K \dot{-} \phi$ the *gentle* epistemic entrenchment contraction of K by ϕ. We shall presently see that severe and gentle epistemic entrenchment contractions may serve, in certain contexts, as upper and lower bounds for Lindström and Rabinowicz's interpolation thesis.

Clearly, the logical properties of the theory contraction operators $\dot{-}$ and $\ddot{-}$ depend on the properties of the entrenchment relation $<$. For information about this dependency, see again Gärdenfors and Makinson (1988) and Rott (1992b, 1996).

9 Theory contraction operations and Simple Filtering

We are now going to investigate the contraction functions that result from applying the entrenchment relations we have generated from a belief base $\mathcal{H}$ on the level of the logically closed set $K = Cn(H)$.

It is easy to show that the entrenchment relation $<^+$ is not suitable for the purposes in question if we use either one of the recipes of Section 8.

Example 3. A base $H = \{p, q\}$ with empty $\prec$, say, intuitively suggests that $K \dot{-} p$ should contain q, since q is not in K just because p is in H or K. However, we have neither $p <^+ q$ nor $p <^+ p \vee q$. The severe entrenchment contraction based on $<^+$ gives us $K \dot{-} p = Cn(\emptyset)$ while the gentle one gives us $K \dot{-} p = Cn(p \supset q)$. Neither result is intuitively acceptable.

It should be noted that the defect of entrenchment contractions based on $<^+$ does not consist in a violation of the filtering condition, but rather of a principle of *Converse Filtering*:

(CF) The contraction of one's beliefs by ϕ should contain all those sentences that were believed, but not believed "just because" ϕ was believed.

In contrast to (SF) and (F), (CF) recommends a sort of doxastic conservatism or minimal change. A special case of failure to meet this principle is the "drowning effect" that was identified by Benferhat et al. (1993) and will be illustrated in the next section.

Now let us work with the negative entrenchment relation $<^-$ constructed from $\mathcal{H}$. Which definition of contractions based on $<^-$ is the better one, (Def$\dot{-}$) or (Def$\ddot{-}$)? It is hard to give a categorical answer. Lindström and Rabinowicz's (1991) abstain from recommending either the gentle or the severe contraction recipe. They argue that these extremes should be taken as upper and lower bounds and that any contraction function lying "between" them is acceptable. Let us call this suggestion Lindström and Rabinowicz's *interpolation thesis*. In Rott (1992c, pp. 48–49), it is shown how their thesis (as well as their idea of an entrenchment relation with incomparabilities) becomes relevant for the reconstruction of belief base contractions at the "knowledge level":

Observation 4. Let $\dot{-}$ be the prioritized base contraction function determined by the prioritized belief base $\mathcal{H}$, and let $\ddot{-}$ be the severe and $\dot{-}'$ be the gentle EE-contraction with respect to $<^-$. Then

$$K\ddot{-}\phi \quad \subseteq \quad K\dot{-}\phi \quad \subseteq \quad K\dot{-}'\phi$$

The converse inclusions are not valid, even when the priority relation $\prec$ is empty.

I take it that our intuitions are favourable to prioritized base contractions. The following examples indicate that one cannot get a closer approximation on the theory level than characterized by these bounds.

Example 4. Consider the unprioritized belief base

$$H = \{p, q\}$$

Here q is in $K = Cn(H)$ not just because p is in K, and therefore we expect $K\dot{-}p$ to be $Cn(q)$. On the other hand, in the belief base

$$H' = \{p \wedge q, p \vee q\}$$

q is in $K' = Cn(H')$ just because p is in K', and we expect $K\dot{-}p$ to be $Cn(p \vee q)$. Unfortunately, our relation $<^-$ does not reflect this difference between H and H'. In the case of H, the unique best way to discard p is the set $\{q\}$, and in the case of H' it is $\{p \vee q\}$. In both bases, then, a sentence ϕ is $<^-$-better than p if and only if the unique best way to discard ϕ is the empty set, $\emptyset$. So in both cases, ϕ is $<^-$-better than p if and only if ϕ is in $Cn(p \vee q)$ (assuming classical reasoning with disjunctions):

$$p <^- \phi \quad \text{if and only if} \quad p \vee q \vdash \phi$$

Considering the intuitive difference in the results of contracting the bases H and H' by p, it appears that (Def$\dot{-}$) is the appropriate definition in the case of H, whereas (Def$\ddot{-}$) is the appropriate definition in the case of H'.

We note that the relation $<^-$ is exactly the same for H and H'. For instance, it is easy to calculate that for both H and H'

$$\phi <^- p \vee q \quad \text{if and only if} \quad p \vee q \nvdash \phi$$

The negative entrenchment relation generated by a belief base fails to mirror, in a one-to-one fashion, the dependencies encoded in the different syntactical structures of the belief bases.

Finally, we can show that gentle and severe entrenchment contractions based on $<^-$ satisfy versions of the condition of Simple Filtering.

Observation 5. Given some prioritized belief base $\mathcal{H}$, the theory contraction function $\dot{-}$ over $K = Cn(H)$ defined by (Def$\dot{-}$) using $<^-$ satisfies the Simple Filtering condition in the versions (BSF) and (BJB) for basic beliefs, but not in the versions (DSF) and (DJB) for derived beliefs.

Proof. (BSF) with (BJB) is satisfied. Let ψ be in K just because ϕ is in H, that is, ψ is in $Cn(H)$ but not in $Cn(H - \{\phi\})$. We have to show that ψ is not in $K \dot{-} \phi$, that is, by (Def$\dot{-}$), that $\phi \not<^- \phi \vee \psi$, that is, by the definition of $<^-$, that

there is a $G \subseteq H$ such that $G \not\vdash \phi \vee \psi$ and
for every $G' \subseteq H$ such that $G' \not\vdash \phi$ it does not hold that $G \lhd^- G'$

Consider $G = H - \{\phi\}$. First we show that $H - \{\phi\} \not\vdash \phi \vee \psi$. Suppose for reductio that $H - \{\phi\} \vdash \phi \vee \psi$. Since $H \vdash \psi$, we also have $H - \{\phi\} \vdash \phi \supset \psi$ (assuming that Cn satisfies the deduction theorem). So $H - \{\phi\} \vdash \psi$, contradicting the fact that ψ is in K just because ϕ is in H. Now take any subset G' of H such that $G' \not\vdash \phi$. Then of course ϕ is not in G'. But then $G' - (H - \{\phi\})$ is empty, and hence, by the definition of $\lhd^-$, it cannot hold that $H - \{\phi\} \lhd^- G'$.

The following counterexample shows that (DSF) with (DJB) is not satisfied. Consider the (unprioritized) base $H = \{p \vee q, p \wedge q\}$. Clearly, q is in $K = Cn(H)$ just because p is in K. On the other hand, the unique $\lhd^-$-best way of retracting $p \vee q$ is $\emptyset$, and the unique $\lhd^-$-best way of retracting p is $\{p \vee q\}$. So $p <^- p \vee q$. So, according to (Def$\dot{-}$) using $<^-$, q is in $K \dot{-} p$, in violation to (DSF) with (DJB). QED

Observation 6. Given some prioritized belief base $\mathcal{H}$, the theory contraction function $\ddot{-}$ over $K = Cn(H)$ defined by (Def$\ddot{-}$) using $<^-$ satisfies the Simple Filtering condition (DSF) with (DJB).

Proof. (DSF) with (DJB) is satisfied. Let ψ be in K just because ϕ is in K, that is, ψ is in $Cn(H)$ but not in $Cn(G)$, for all subsets G of H with $G \not\vdash \phi$. We have to show that ψ is not in $K \ddot{-} \phi$, that is, by (Def$\ddot{-}$), that $\phi \not<^- \psi$, that is, by the definition of $<^-$, that

there is a $G \subseteq H$ such that $G \not\vdash \psi$ and
for every $G' \subseteq H$ such that $G' \not\vdash \phi$ it does not hold that $G \lhd^- G'$

Take an arbitrary $\lhd^-$-best $G \subseteq H$ such that $G \not\vdash \psi$. Clearly, such a set G exists. Suppose for reductio that there is a $G' \subseteq H$ such that $G' \not\vdash \phi$ and $G \lhd^- G'$. Since ψ is in K just because ϕ is in K, we get that $G' \not\vdash \psi$. But then, since $G \lhd^- G'$, G is not a $\lhd^-$-best subset of H that does not imply ψ, contradicting our choice of G. QED

Since (DSF) has been given preference over (BSF), it may seem that Observations 5 and 6 taken together provide an argument in favour of severe contractions (and against gentle contractions). However, such a diagnosis cannot be made without qualification. First, it may be doubted that (DJB) presents a plausible analysis of the phrase "just because." And secondly, the Filtering Principle (SF) represents only one side of our intuitions about rational base contractions. It has to be complemented by the equally legitimate intuition expressed by (CF) which – in contrast to (SF) – encourages doxastic conservatism. In this respect, of course, gentle contractions fare better than severe ones.[11]

[11] Compare Pagnucco and Rott (1998).

10 An ordinal base revision mechanism and its problems

Mary-Anne Williams (1995) introduced the concept of a partial entrenchment ranking. Such a ranking is a function $\mathbf{E}$ that assigns to each sentence of a finite base a non-negative integer. A partial entrenchment ranking corresponds to a prioritized belief base $\mathcal{H}$ in which some of the H_i's may be empty. In addition, Williams's entrenchment rankings must satisfy some logical constraints which are not important for our discussion, however.

Our above function ent^+ may still be used in this context, with a slight change of definition.[12] Given a partial entrenchment ranking $\mathbf{E}$, the entrenchment of an arbitrary sentence ϕ which need not be in the domain of $\mathbf{E}$, in symbols $\text{ent}^+(\phi)$, is the maximal number k such that $\phi \in Cn(\{\psi \in \text{dom}(\mathbf{E})\colon \mathbf{E}(\psi) \geq k\})$. Equivalently, $\text{ent}^+(\phi) = \max\{\text{qual}^+(G)\colon G \subseteq H \text{ and } G \vdash \phi\}$, with $\text{qual}^+(G)$ defined as $\min\{\mathbf{E}(\psi)\colon \psi \in G\}$. If ϕ is not in $Cn(H)$, we put $\text{ent}^+(\phi) = 0$ (i.e., $\max(\emptyset) = 0$). It is tacitly understood in what follows that ent^+ is constructed in this way from a given entrenchment ranking $\mathbf{E}$. A sentence ϕ is believed if and only if $\text{ent}^+(\phi) > 0$.

Williams addresses the problem of how to revise partial entrenchment rankings in the face of new information. She proposes an ordinal version of the usual belief change operations, one that allows an input sentence ϕ to come together with an integer i specifying the *a posteriori* entrenchment of ϕ. As pointed out already by Spohn (1988), one of the advantages of ordinalized belief representations is that they provide for modellings of iterations of belief changes – something that is much harder to model if there is no ordinal structure to rely upon. A disadvantage of ordinalized representations is that they presuppose that all beliefs are comparable in terms of epistemic entrenchment. In the terminology introduced above, Williams's account is tied to the positive concept of epistemic entrenchment.

Williams (1995) suggests the following method of revising partial entrenchment rankings:

$$\mathbf{E}^*(\phi, i) = \begin{cases} \mathbf{E}^-(\phi, i) & \text{if } i < \text{ent}^+(\phi) \\ (\mathbf{E}^-(\neg\phi, 0))^+(\phi, i) & \text{otherwise} \end{cases}$$

where use is made of the contraction method

$$\mathbf{E}^-(\phi, i)(\psi) = \begin{cases} i & \text{if } \text{ent}^+(\phi) = \text{ent}^+(\phi \vee \psi) \\ & \quad \text{and } \mathbf{E}(\psi) > i \\ \mathbf{E}(\psi) & \text{otherwise} \end{cases}$$

for all $\psi \in \text{dom}(\mathbf{E})$, and the expansion method

$$\mathbf{E}^+(\phi, i)(\psi) = \begin{cases} \mathbf{E}(\psi) & \text{if } \mathbf{E}(\psi) > i \\ i & \text{if } \phi \equiv \psi \\ & \quad \text{or } \mathbf{E}(\psi) \leq i < \text{ent}^+(\neg\phi \vee \psi) \\ \text{ent}^+(\neg\phi \vee \psi) & \text{otherwise} \end{cases}$$

[12] $\text{ent}^+(\phi)$ is called the degree of acceptance of ϕ by Williams.

for all $\psi \in \mathrm{dom}(\mathbf{E} \cup \{\phi\})$.[13]

Unfortunately, this suggestion yields counterintuitive results. The following example taken from Williams (1995) shows that it suffers from what has been called the "drowning effect" by Benferhat et al. (1993).

Example 5. Consider the entrenchment ranking $\mathbf{E}$ over the set $H = \{p, p \supset q, r\}$ with the values $\mathbf{E}(p) = 1$, $\mathbf{E}(p \supset q) = 3$ and $\mathbf{E}(r) = 2$. Williams discusses only changes at the lowest level, namely changes with respect to p which involve no interference with higher levels. At this level, her revision methods work fine. Problems arise, however, if we look at changes at higher levels. Consider for instance a contraction with respect to r, which gives us the equation

$$\mathbf{E}^-(r,0)(\psi) = \begin{cases} 0 & \text{if } \mathrm{ent}^+(r) = \mathrm{ent}^+(r \vee \psi) \\ \mathbf{E}(\psi) & \text{otherwise} \end{cases}$$

for each ψ in $\mathrm{dom}(\mathbf{E})$. Since $\mathrm{ent}^+(r) = 2 = \mathrm{ent}^+(r \vee p)$, we get $\mathbf{E}^-(r,0)(p) = 0$. That is, p is no longer believed after a contraction with respect to r. This method does not do justice to the belief base because p is a basic belief which is independent of r.

We conclude that Williams's contraction and revision operations gratuitously lose independent beliefs with a low priority in a belief base. Similar problems arise for the base change mechanisms in Williams (1994, pp. 96–100) and Wobcke (1995, pp. 78–79).[14]

The above ordinal base contraction mechanism yields wrong results, because the transition from $\mathbf{E}$ to ent^+ in effect implements the positive concept of epistemic entrenchment, and we have already seen that this concept is not suitable for use in belief base dynamics. Although, in the example, the basic belief r is independent of the other basic belief p, the entrenchment of $p \vee r$ is not higher than that of r alone. For this reason the definition of $\mathbf{E}^-(\phi, i)$ (which in effect implements an ordinalized version of gentle epistemic entrenchment contractions) cannot sensibly be applied.

This kind of "drowning effect" would be avoided, if we could use the finer tuned (but not connected) negative entrenchment relation $<^-$. If we apply the negative concept of entrenchment in Williams's example, we find that $r <^- p \vee r$, since $H \perp r$ is $\{\{p \supset q, p\}\}$ and $H \perp p \vee r$ is $\{\{p \supset q\}\}$. This relation takes the elements of a belief base more seriously as independent items of information.

11 Conclusion

We have formalized several ways of accounting, in the context of logically closed theories, for foundationalist intuitions that underlie change operations applying

[13] In the definition of $\mathbf{E}^+$, it is not quite clear which line is applicable when both $\phi \equiv \psi$ and $\mathbf{E}(\psi) > i$.

[14] Wobcke in addition violates Maxim B. In more recent work, Williams (1996) has revised her approach so that the drowning effect is avoided.

to belief bases. A positive and a negative concept of entrenchment was defined on the basis of the structure of a given, possibly prioritized belief base. Only the latter, more fine-grained concept proved to be appropriate for a successful attempt at approximating base changes on the theory level. We investigated the question to which degree we can comply with the fundamental intuition expressed by the various Filtering Conditions which say that all (and only) beliefs that are believed "just because" a retracted belief was believed should be withdrawn.

Acknowledgement

I would like to thank Andreas Herzig, Emil Weydert and Mary-Anne Williams for a number of stimulating disputations. – An earlier version of this paper has appeared in S.O. Hansson, W. Rabinowicz and K. Segerberg (eds.), *Logic for a Change. Essays Dedicated to Sten Lindström on the Occasion of his 50th Birthday*, Uppsala Prints and Preprints in Philosophy 1995:9, pp. 106–124.

References

Anderson, Alan Ross, Nuel D. Belnap and J. Michael Dunn (1992): *Entailment. The Logic of Relevance and Necessity*, Vol. II, Princeton University Press, Princeton.

Belnap, Nuel D. (1979): 'Rescher's Hypothetical Reasoning: An Amendment', in E. Sosa (ed.), *The Philosophy of Nicholas Rescher*, Reidel, Dordrecht 1979, pp. 19–28.

Benferhat, Salem, Claudette Cayrol, Didier Dubois, Jérôme Lang and Henri Prade (1993): 'Inconsistency Management and Prioritized Syntax-Based Entailment', in R. Bajcsy (*ed.*), *Proceedings of the 13th International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann, pp. 640–645.

Cayrol, Claudette, and Marie-Christine Lagasquie-Schiex (1995): 'Non-monotonic Syntax-Based Entailment: A Classification of Consequence Relations', in C. Froidevaux and J. Kohlas (eds.), *Symbolic and Quantitative Approaches to Reasoning and Uncertainty – ECSQARU '95*, LNCS **946**, Berlin: Springer, pp. 107–114.

del Val, Alvaro (1997): 'Non-Monotonic Reasoning and Belief Revision: Syntactic, Semantic, Foundational, and Coherence Approaches', *Journal of Applied Non-Classical Logics* **7**, 213–240.

Doyle, Jon (1992): 'Reason Maintenance and Belief Revision: Foundations vs. Coherence Theories', in P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, Cambridge, pp. 29–51.

Dubois, Didier, and Henri Prade (1991): 'Epistemic Entrenchment and Possibilistic Logic', *Artificial Intelligence* **50**, 223–239.

Dubois, Didier, and Henri Prade (1992): 'Belief Change and Possibility Theory', in P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, pp. 142–182.

Fermé, Eduardo, and Ricardo Rodriguez (1998): 'A Brief Note About the Rott Contraction', to appear in the *Logic Journal of the IGPL*.

Fuhrmann, André (1991): 'Theory Contraction Through Base Contraction', *Journal of Philosophical Logic* **20**, 175–203.

Gärdenfors, Peter (1990): 'The Dynamics of Belief Systems: Foundations vs. Coherence Theories', *Revue Internationale de Philosophie* **44**, 24–46.

Gärdenfors, Peter, and David Makinson (1988): 'Revisions of Knowledge Systems Using Epistemic Entrenchment', in M. Vardi (*ed.*), *Theoretical Aspects of Reasoning About Knowledge*. Los Altos, CA: Morgan Kaufmann, pp. 83–95.

Hansson, Sven Ove (1993): 'Changes on Disjunctively Closed Bases', *Journal of Logic, Language and Information* **2**, 255–284.

Hansson, Sven Ove (1994): 'Taking Belief Bases Seriously', in D. Prawitz and D. Westerståhl (*eds*), *Logic and Philosophy in Uppsala*. Synthese Library, Volume 236, Dordrecht: Kluwer Academic Publishers, pp. 13–28.

Jaskowski, Stanislaw (1948/1969): 'Propositional Calculus for Contradictory Deductive Systems', *Studia Logica*, Vol. 24 (1969), pp. 143–157; first published as 'Rachunek zdah dla systemow dedukcyjnych sprzecznych', *Studia Societatis Scientiarun Torunesis, Sectio A*, Vol. I (1948), No. 5, pp. 55–77.

Kyburg, Henry E. Jr. (1970): 'Conjunctivitis', in M. Swain (ed.), *Induction, Acceptance, and Rational Belief*, Reidel, Dordrecht, pp. 55-82. (Reprinted in H.E.K. Jr., *Epistemology and Inference*, University of Minnesota Press, Minneapolis 1983, pp. 232–254.)

Kyburg, Henry E. Jr. (1997): 'The Rule of Adjunction and Reasonable Inference', *Journal of Philosophy* 94, 109–125.

Lindström, Sten, and Wlodzimierz Rabinowicz (1991): 'Epistemic Entrenchment with Incomparabilities and Relational Belief Revision', in A. Fuhrmann and M. Morreau (eds.), *The Logic of Theory Change*, LNCS **465**, Berlin: Springer, pp. 93–126.

Martins, Joao P., and Stuart C. Shapiro (1988): 'A Model for Belief Revision', *Artificial Intelligence* **35**, 25–79.

Nebel, Bernhard (1989): 'A Knowledge Level Analysis of Belief Revision', in R. Brachman, H. Levesque and R. Reiter (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference*, San Mateo, Ca., Morgan Kaufmann, pp. 301–311.

Olsson, Erik J. (1997): *Coherence – Studies in Epistemology and Belief Revision*, Acta Universitatis Upsaliensis, Uppsala.

Pagnucco, Maurice, and Hans Rott (1998): 'Severe Withdrawal – and Recovery', to appear in the *Journal of Philosophical Logic*.

Rao, Anand S., and Norman Y. Foo (1989): 'Formal Theories of Belief Revision', in R. Brachman, H. Levesque, and R. Reiter, (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference*, San Mateo, Ca., Morgan Kaufmann, pp. 369–380.

Rescher, Nicholas (1964): *Hypothetical Reasoning*, North-Holland, Amsterdam.

Rescher, Nicholas (1976): *Plausible Reasoning*, Van Gorcum, Assen, Amsterdam.

Rescher, Nicholas (1979): 'Reply to Belnap', in E. Sosa (ed.), *The Philosophy of Nicholas Rescher*, Reidel, Dordrecht, pp. 29–31.

Rescher, Nicholas, and Robert Brandom (1979), *The Logic of Inconsistency: A Study in Nonstandard Possible-World Semantics and Ontology*, Basil Blackwell, Oxford.

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

Rott, Hans (1991b): 'A Non-monotonic Conditional Logic for Belief Revision I', in A. Fuhrmann and M. Morreau (eds.), *The Logic of Theory Change*. LNCS **465**, Berlin: Springer, pp. 135–81.

Rott, Hans (1992a) 'On the Logic of Theory Change: More Maps Between Different Kinds of Contraction Function', in P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, Cambridge, pp. 122–141.

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

Rott, Hans (1992c): 'Modellings for Belief Change: Prioritization and Entrenchment', *Theoria* **58**, pp. 21–57.

Rott, Hans (1996): *Making Up One's Mind. Foundations, Coherence, Nonmonotonicity*, Habilitationsschrift, University of Konstanz.

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

Urchs, Max (1995): 'Discursive Logic', *Studia Logica* **54**, 231–249.

Williams, Mary-Anne (1994): 'On the Logic of Theory Base Change', in C. MacNish, D. Pearce and L.M. Pereira (eds.), *Logics in Artificial Intelligence*, LNCS **838**, Springer, Berlin, pp. 86–105.

Williams, Mary-Anne (1995): 'Iterated Theory Base Change: A Computational Model', in *IJCAI-95 – Proceedings of the 14th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo.

Williams, Mary-Anne (1996): 'Towards a Practical Approach to Belief Revision: Reason-Based Change', in L. Carlucci Aiello and C. Shapiro (eds), *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann Publishers, pp. 412–421.

Williams, Bernard (1970): 'Deciding to Believe', in H.E. Kiefer and M.K. Munitz (eds.), *Language, Belief, and Metaphysics*, State University of New York Press, Albany, pp. 95–111.

Wobcke, Wayne (1995): 'Belief Revision, Conditional Logic and Nonmonotonic Reasoning', *Notre Dame Journal of Formal Logic* 36, 55–102.

Version of October 30, 1998.

This article was processed using the LaTeX macro package with LLNCS style