The lexicographic closure as a revision process

Richard Booth*
University of Leipzig
Dept. of Computer Science
Augustusplatz 10/11
04109 Leipzig, Germany
booth@informatik.uni-leipzig.de

ABSTRACT. *The connections between nonmonotonic reasoning and belief revision are well-known. A central problem in the area of nonmonotonic reasoning is the problem of default entailment, i.e., when should an item of default information representing "if θ is true then, normally, ϕ is true" be said to follow from a given set of items of such information. Many answers to this question have been proposed but, surprisingly, virtually none have attempted any explicit connection to belief revision. The aim of this paper is to give an example of how such a connection can be made by showing how the lexicographic closure of a set of defaults may be conceptualised as a process of iterated revision by sets of sentences. Specifically we use the revision method of Nayak.*
KEYWORDS: Belief revision, default reasoning, iterated revision, revision by sets, epistemic entrenchment, rational consequence.

Introduction and preliminaries

The methodological connections between the areas of nonmonotonic reasoning, i.e., the process by which an agent may, possibly, withdraw previously derived conclusions upon enlarging her set of hypotheses ([Mak 94]), and belief revision, i.e., the process by which an agent changes her beliefs upon discovering some new information ([AGM 85, Gär 88]), are well-known (see, for example, [GM 94, GR 95, MG 90, Rot 96]). As a consequence, it is possible to translate particular problems in one area into problems in the other. One particular problem in nonmonotonic reasoning is the question of *default entailment*, i.e., when should we regard one item of so-called "default knowledge" (hereafter just "default"), i.e., an expression of the form $\theta \Rightarrow \phi$ standing for "if θ then normally (or usually, or typically) ϕ", as "following from" a given set of defaults. Several answers to this question have been proposed in the literature (such as in [BCDLP 93, BSS 95, Bre 89, GMP 93, Leh 95, LM 92, Pea 90, Wey 96], to

*Supported by the DFG project "Computationale Dialektik" within the DFG research group "Kommunikatives Verstehen".

name but a few) but none of them (with the exception of [Wey 96], but see also Section 5 of this paper) seem to attempt any explicit connection with belief revision. The aim of this paper is to make a start on such a connection by showing how one particular method of default entailment, namely the lexicographic closure construction ([BCDLP 93, Leh 95]) can be given a formulation in terms of a certain method of belief revision which was first given by Nayak [Nay 94] and studied further by Nayak, Nelson and Polansky [NNP 96]. In the process, we shall uncover one or two interesting avenues for further research on both sides.

The plan of this paper is as follows. In Section 1 we formally pose the basic question of default entailment outlined above and describe the lexicographic closure. The set of defaults defined by the lexicographic closure, considered as a binary relation, forms a rational consequence relation (in the sense of [KLM 90]). This means that it may be described by a sequence of sets of worlds of the underlying logical language, equivalently a total pre-order on the set of worlds. Section 2 introduces the theory of belief revision and the important notion of epistemic entrenchment relation (E-relation for short) which it utilises. Also in this section we describe in detail the correspondence between E-relations and rational consequence relations and show that, in effect, they are two different ways of describing the same thing. In particular, epistemic entrenchment relations may also be represented as sequences of world-sets. We use this method of representation to describe Nayak's operation of revision in Section 3. Nayak proposes to model revision of an epistemic state (represented as an E-relation) by an arbitrary set of sentences by first converting this set into an E-relation and then revising by this relation. In Section 4 we present one particular method, the idea behind which crops up several places in the literature, for generating an E-relation from a set of sentences and show our main result: that, given this method, the E-relation corresponding to the lexicographic closure can be obtained by revising the maximally ignorant epistemic state (i.e., the E-relation in which the only sentences believed are the tautologies) firstly by the set of (the material counterparts of) those defaults which are the least specific (according to an accepted relation of *is more specific than* among defaults), then those defaults which are the next-least specific and so on up to the set of the most specific defaults. In Section 5 we take a brief look at another method of default entailment, due to Brewka, which is related to the lexicographic closure and which has already been shown to have very close connections with a different type of belief revision operator due to Nebel. Finally, in Section 6 we give our ideas for possible further study before offering some short concluding remarks.

Before we get started, let us fix our notation. Throughout this paper, L is an arbitrary but fixed propositional language built up from a finite set of propositional variables using the usual connectives $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$ and constants $\top, \bot$. Semantics is provided by the (finite) set W of propositional worlds. For $w \in W$ and $\theta \in L$ we write $w \models \theta$ whenever w satisfies θ, and set $S_\theta = \{w \in W \mid w \models \theta\}$. Given $E \cup \{\phi\} \subseteq L$ we write $E \models \phi$ whenever $\bigcap_{\theta \in E} S_\theta \subseteq S_\phi$ and let $Cn(E)$ denote the set $\{\phi \mid E \models \phi\}$. We say E is *inconsistent* if $E \models \bot$, otherwise E is *consistent*. As usual we write $\theta \models \phi$ rather than $\{\theta\} \models \phi$, write $\models \theta$ rather than $\emptyset \models \theta$, and say "θ is inconsistent" rather than "$\{\theta\}$ is inconsistent". For any $w \in W$ and $E \subseteq L$ we set $sent_E(w) = \{\theta \in E \mid w \models \theta\}$, i.e., $sent_E(w)$ is the set of sentences in E which w satisfies, while, for finite $E \subseteq L$, $\bigwedge E$, respectively $\bigvee E$, denotes the conjunction, respectively disjunction, in some order, of all the sentences in E. Since for our purposes the precise order used here will always be irrelevant, we leave it unspecified. In some of our proofs we will often be seen to treat worlds in W as if they are sentences, for example writing $\neg w$ or $\bigvee \mathcal{U}$ for $\mathcal{U} \subseteq W$. Whenever a world w appears in the scope of a connective like this it is to be understood that we are using w to represent any sentence $\alpha \in L$ such that $S_\alpha = \{w\}$. (Again the precise choice here will be irrelevant.) Such a sentence can be found by conjoining, in some order, all the propositional variables in L which w satisfies with all the negations of the variables which w does not satisfy. Note that this gives us $\mathcal{U} = S_{\bigvee \mathcal{U}}$ for any $\mathcal{U} \subseteq W$. Finally, for an arbitrary set X we use $|X|$ to denote the cardinality of X.

1 The lexicographic closure of a set of defaults

Suppose we have somehow learnt that an intelligent, rational agent believes some finite set of defaults $\Delta = \{\theta_i \Rightarrow \phi_i \mid \theta_i, \phi_i \in L, i = 1, \ldots, l\}$. In this case what other assertions of this form should we conclude our agent believes? Or, put another way, what is the binary relation $\mathrel{|\!\sim}^\Delta$ on L where $\theta \mathrel{|\!\sim}^\Delta \phi$ holds iff we can conclude, on the basis of Δ, that if θ is true then, normally, ϕ will also be true? In order to be able to answer this question we need some criteria by which we can judge possible answers. These criteria fall into two different categories. Firstly we need *global* criteria which are concerned with the mapping $\Delta \mapsto \mathrel{|\!\sim}^\Delta$ (e.g., a minimal requirement would seem to be $\theta \Rightarrow \phi \in \Delta$ implies $\theta \mathrel{|\!\sim}^\Delta \phi$). Secondly we need *local* criteria which are simply concerned with the internal closure properties of $\mathrel{|\!\sim}^\Delta$ (e.g., $\mathrel{|\!\sim}^\Delta$ should be *reflexive*, i.e., $\theta \mathrel{|\!\sim}^\Delta \theta$ for all θ). In this paper, one answer which we are particularly interested in is the lexicographic closure construction which was proposed independently by both Benferhat et al [BCDLP 93] and Lehmann [Leh 95]. We describe this construction now.

Throughout this paper we assume that Δ is an arbitrary but fixed, finite set of defaults. For this paper we also make the simplifying assumption that Δ is "consistent", in the sense that its set of *material counterparts* $\Delta^\rightarrow \subseteq L$ is consistent, where, for an arbitrary set of defaults Σ, we set $\Sigma^\rightarrow = \{\lambda \rightarrow \chi \mid \lambda \Rightarrow \chi \in \Sigma\}$. Using a procedure given in [Pea 90] (or, equivalently, in [Leh 95]) we may partition Δ into $\Delta = (\Delta_0, \ldots, \Delta_n)$, where the Δ_i correspond, in a precise sense, to "levels of specificity" – given a default $\delta \in \Delta$, the larger the i for which $\delta \in \Delta_i$, the more specific are the situations to which δ is applicable.

Following [Pea 90], we call this partition the *Z-partition* of Δ.[1] Like many methods of default entailment (see [BCDLP 93] for several examples, one of which we reproduce in Section 5), the lexicographic closure can be based on a method of choosing maximal consistent subsets of $\Delta^{\rightarrow}$. More precisely the lexicographic closure is a member of a family of consequence relations $\mid\!\sim^{\Delta}_{\ll}$, where $\ll$ is an ordering on $2^{\Delta^{\rightarrow}}$, and, for all $\theta, \phi \in L$, we have

$$\theta \mid\!\sim^{\Delta}_{\ll} \phi \quad \text{iff} \quad \begin{array}{l}\text{for all } \Gamma \subseteq \Delta^{\rightarrow} \text{ such that } \Gamma \cup \{\theta\} \text{ is consistent and } \Gamma \text{ is} \\ \ll\text{-maximal amongst such subsets, we have } \Gamma \cup \{\theta\} \models \phi,\end{array}$$

To specify the lexicographic closure, we instantiate the order $\ll$ above, with the help of the Z-partition, as follows: Given subsets $A, B \subseteq \Delta^{\rightarrow}$ let $A_i = A \cap \Delta_i^{\rightarrow}$ and $B_i = B \cap \Delta_i^{\rightarrow}$ for each $i = 0, \ldots, n$. We define an ordering $\ll_{lex}$ on $2^{\Delta^{\rightarrow}}$ by:

$$A \ll_{lex} B \quad \text{iff} \quad \begin{array}{l}\text{there exists } i \text{ such that } |A_i| < |B_i| \text{ and,} \\ \text{for all } j > i,\ |A_j| = |B_j|.\end{array}$$

(The reason for the name "lexicographic closure" should now be clear.) The lexicographic closure $\mid\!\sim^{\Delta}_{lex}$ is then just defined to be $\mid\!\sim^{\Delta}_{\ll_{lex}}$.

How successful is $\mid\!\sim^{\Delta}_{lex}$ in achieving the goals of default reasoning? We refer the reader to [Leh 95] for the details concerning its global properties. The local properties of $\mid\!\sim^{\Delta}_{lex}$ are summed up by the following proposition, which can be found jointly in [BCDLP 93] and [Leh 95].

Proposition 1 *The binary relation $\mid\!\sim^{\Delta}_{lex}$ is a rational consequence relation, i.e., it satisfies the following properties of a binary relation $\mid\!\sim$ on L: For all $\theta, \phi, \psi \in L$,*

$\theta \mid\!\sim \theta$ *(Reflexivity)*

$$\frac{\theta \mid\!\sim \phi,\ \models \theta \leftrightarrow \psi}{\psi \mid\!\sim \phi} \ \textit{(Left Logical Equivalence)}$$

$$\frac{\theta \mid\!\sim \phi,\ \phi \models \psi}{\theta \mid\!\sim \psi} \ \textit{(Right Weakening)}$$

$$\frac{\theta \mid\!\sim \phi,\ \theta \mid\!\sim \psi}{\theta \mid\!\sim \phi \wedge \psi} \ \textit{(And)}$$

[1] More precisely, the Z-partition $(\Delta_0, \ldots, \Delta_n)$ can be defined inductively by setting

$$\Delta_0 = \{\theta \Rightarrow \phi \in \Delta \mid \Delta^{\rightarrow} \not\models \neg\theta\}$$

and, for $i \geq 1$,

$$\Delta_i = \{\theta \Rightarrow \phi \in (\Delta - \bigcup_{j<i} \Delta_j) \mid (\Delta^{\rightarrow} - \bigcup_{j<i} \Delta_j^{\rightarrow}) \not\models \neg\theta\}.$$

This process stops as soon as we reach some i for which $\Delta_i = \emptyset$. At this point, letting Δ_∞ denote $\Delta - \bigcup_{j<i} \Delta_j$, we set the Z-partition to be $(\Delta_0, \ldots, \Delta_{i-1})$ if $\Delta_\infty = \emptyset$, or $(\Delta_0, \ldots, \Delta_{i-1}, \Delta_\infty)$ if $\Delta_\infty \neq \emptyset$. For more details we refer the reader to [Pea 90] and [Leh 95].

$$\frac{\theta \mid\sim \phi,\ \psi \mid\sim \phi}{\theta \vee \psi \mid\sim \phi} \textit{ (Or)}$$

$$\frac{\theta \mid\sim \phi,\ \theta \mid\sim \psi}{\theta \wedge \phi \mid\sim \psi} \textit{ (Cautious Monotonicity)}$$

$$\frac{\theta \mid\sim \phi,\ \theta \not\mid\sim \neg\psi}{\theta \wedge \psi \mid\sim \phi} \textit{ (Rational Monotonicity)}$$

Furthermore, $\mid\sim^{\Delta}_{lex}$ *is consistency preserving, i.e., it satisfies*

$$\frac{\theta \mid\sim \bot}{\theta \models \bot} \textit{ (Consistency Preservation)}$$

Hence $\mid\sim^{\Delta}_{lex}$ satisfies thesis 1.1 of [LM 92] which requires the relation $\mid\sim^{\Delta}$ generated by Δ to be a rational consequence relation. Now we already know (see, for example [Fre 93, LM 92]) that rational consequence relations may be represented by finite sequences $\vec{\mathcal{U}} = (\mathcal{U}_0, \ldots, \mathcal{U}_k)$ of mutually disjoint subsets of W in the following sense: Given such a sequence $\vec{\mathcal{U}}$, if we define a binary relation $\mid\sim_{\vec{\mathcal{U}}}$ on L by setting

$$\theta \mid\sim_{\vec{\mathcal{U}}} \phi \quad \text{iff} \quad \begin{array}{ll} \textit{either} & \mathcal{U}_i \cap S_\theta = \emptyset \text{ for all } i = 0, 1, \ldots, k \\ \textit{or} & \mathcal{U}_i \cap S_\theta \subseteq S_\phi \text{ for the least } i \text{ such that } \mathcal{U}_i \cap S_\theta \neq \emptyset \end{array}$$

then $\mid\sim_{\vec{\mathcal{U}}}$ forms a rational consequence relation, while moreover every rational consequence relation arises in this way from some sequence $\vec{\mathcal{U}}$.[2] The intuition behind the sequences $\vec{\mathcal{U}}$ is that they represent a "ranking" of the worlds in W according to their plausibility – the lower the i for which $w \in \mathcal{U}_i$, the more plausible, or less exceptional, in relation to the other worlds, it is considered to be. If $w \notin \mathcal{U}_i$ for all i then we may take w to be considered "impossible". The definition of $\mid\sim_{\vec{\mathcal{U}}}$ can then be translated as saying that, given θ, we should conclude that ϕ is normally true iff either θ is considered impossible or all the most plausible worlds which satisfy θ also satisfy ϕ. It should be kept in mind that, although here we are using simple non-negative integers to index the elements in $\vec{\mathcal{U}}$, there is in general nothing to stop us from using members of *any* linearly ordered set.

One thing to note about the definition of $\mid\sim_{\vec{\mathcal{U}}}$ given above is that we are allowing $\emptyset$ to appear, possibly more than once, in $\vec{\mathcal{U}}$.[3] This freedom comes in useful when proving some of our results. It also has the effect that the mapping $\vec{\mathcal{U}} \mapsto \mid\sim_{\vec{\mathcal{U}}}$ detailed above is not injective – given a rational consequence relation $\mid\sim$ there will be many (in fact infinitely many) sequences $\vec{\mathcal{U}}$ such that

[2] Such sequences are clearly equivalent to the ranked models used to characterise rational consequence relations in [LM 92].

[3] This approach carries us very close to the "semi-quantitative" approaches of [Spo 88, Wey 96, Wil 94], which use an explicit ranking function as a starting point rather than deriving one from a sequence of world-sets. Our approach, though, remains squarely qualitative in character.

$\mid\!\sim = \mid\!\sim_{\vec{\mathcal{U}}}$.[4] Another thing to note about $\mid\!\sim_{\vec{\mathcal{U}}}$ is that, as is easily verified,$\mid\!\sim_{\vec{\mathcal{U}}}$ will be consistency preserving iff $\bigcup_{i=0}^{k} \mathcal{U}_i = W$, i.e., iff all worlds are considered possible, while it will be trivial, i.e., will satisfy $\theta \mid\!\sim_{\vec{\mathcal{U}}} \phi$ for *all* θ and ϕ, iff $\bigcup_{i=0}^{k} \mathcal{U}_i = \emptyset$, i.e., iff *no* worlds are considered possible. We make the following definitions:

Definition 1 *Let $\vec{\mathcal{U}} = (\mathcal{U}_0, \ldots, \mathcal{U}_k)$ be a finite sequence of mutually disjoint subsets of W. We shall say that $\vec{\mathcal{U}}$ is* full *iff $\bigcup_{i=0}^{k} \mathcal{U}_i = W$ and that $\vec{\mathcal{U}}$ is* empty *iff $\bigcup_{i=0}^{k} \mathcal{U}_i = \emptyset$. We let Υ denote the set of all such $\vec{\mathcal{U}}$ which are either full or empty.*

Hence Proposition 1 tells us that there must exist a full sequence $\vec{\mathcal{U}} \in \Upsilon$ such that $\theta \mid\!\sim_{lex}^{\Delta} \phi$ iff $\theta \mid\!\sim_{\vec{\mathcal{U}}} \phi$. What form does $\vec{\mathcal{U}}$ take here? The answer is given in [BCDLP 93] and [Leh 95] (and is, in fact, used to *define* $\mid\!\sim_{lex}^{\Delta}$ in the latter). In this paper we show that we can arrive at this answer via a different route.

2 Belief revision and epistemic entrenchment

Belief revision has been an active area of research, both in philosophy and computer science, since the early 1980's. It is concerned with the following problem: How should an agent revise her beliefs upon receiving some new information which may, possibly, contradict some of her current beliefs? The most popular basic framework within which this question is studied is the one laid down by Alchourŕon, Gärdenfors and Makinson (AGM) in [AGM 85]. In that framework an agent's epistemic state is represented as a deductively closed set of sentences K called a *belief set*, and the new information, or epistemic input, is represented as a single sentence θ. AGM propose a number of postulates which a reasonable operation of revision $K * \theta$ should satisfy. In particular, the revised belief set should contain the epistemic input, i.e., $\theta \in K * \theta$, and should be consistent, i.e., $\perp \notin K * \theta$.[5] In order to meet these requirements, in the general case when the input is inconsistent with the prior belief set, the agent is forced to give up some of her prior beliefs. One way of determining precisely *which* sentences the agent should give up in this situation is to assign to the agent an E-relation $\preceq$ on L (see, for example, [Gär 88, GM 94, Nay 94, Rot 92a, Rot 96]).

The intuitive meaning behind E-relations $\preceq$ is that $\phi \preceq \psi$ should hold iff the agent finds it at least as easy to give up ϕ as she does ψ, i.e., her belief in ψ is at least as entrenched as her belief in ϕ. In cases of conflict the agent should then give up those sentences which are less entrenched. In what follows we use $\prec$ to denote the strict part of $\preceq$, i.e., $\theta \prec \phi$ iff $\theta \preceq \phi$ and not$(\phi \preceq \theta)$. We follow [Nay 94] in formally defining E-relations as follows:

[4] Since clearly we can insert as many copies of $\emptyset$ into the sequence $(\mathcal{U}_0, \ldots, \mathcal{U}_k)$ as we wish without changing the relation $\mid\!\sim_{\vec{\mathcal{U}}}$.

[5] Unless the epistemic input itself is inconsistent. See [Gär 88] for the full list of postulates with detailed discussion.

Definition 2 *An* epistemic entrenchment relation (E-relation) *(on L) is a relation $\preceq \subseteq L \times L$ which satisfies the following conditions for all $\theta, \phi, \psi \in L$,*

(E1)	*If $\theta \preceq \phi$ and $\phi \preceq \psi$ then $\theta \preceq \psi$*	*(transitivity)*
(E2)	*If $\theta \models \phi$ then $\theta \preceq \phi$*	*(dominance)*
(E3)	*$\theta \preceq \theta \wedge \phi$ or $\phi \preceq \theta \wedge \phi$*	*(conjunctiveness)*
(E4)	*Given a $\psi \in L$ such that $\perp \prec \psi$, if $\theta \preceq \phi$ for all $\theta \in L$, then $\models \phi$*	*(maximality)*

Conditions (E1)–(E3) together imply that, for all $\theta, \phi \in L$, we have $\theta \preceq \phi$ or $\phi \preceq \theta$ (so $\theta \prec \phi$ iff not$(\phi \preceq \theta)$). Regarding (E4), it is easy to see that (E1) and (E2) imply that if, for all ψ, we have $\psi \preceq \perp$ then $\theta \preceq \phi$ holds for *all* θ and ϕ, thus in this case $\preceq$ collapses into what we call the *absurd* E-relation. Hence (E4) says that the only sentences which are maximally entrenched are the tautologies, *unless* we are in the special case where $\preceq$ is absurd, in which case *all* sentences – tautologies and contradictions alike – are equally as entrenched as each other. The following derived property of E-relations will be used in the proof of Theorem 1.

Lemma 1 *Let $\preceq$ be an E-relation and let $\theta, \phi \in L$ be such that $\theta \preceq \phi$. Then, for all $\chi \in L$, we have $\chi \preceq \theta$ iff $\chi \preceq \theta \wedge \phi$.*

Proof: Firstly $\chi \preceq \theta \wedge \phi$ implies $\chi \preceq \theta$, without the aid of $\theta \preceq \phi$, by (E1) and (E2). For the converse direction we do need $\theta \preceq \phi$. Suppose $\chi \preceq \theta$. Then, by (E3), either $\theta \preceq \theta \wedge \phi$ or $\phi \preceq \theta \wedge \phi$. If the former holds then $\chi \preceq \theta \wedge \phi$ by (E1), while in the latter case we get $\theta \preceq \theta \wedge \phi$ by (E1) with $\theta \preceq \phi$ and so again $\chi \preceq \theta \wedge \phi$ by (E1). ■

Note that the above definition of E-relation differs from the *standard* definition of E-relations, such as is found in [Gär 88], in two ways. Firstly, in the standard definition, the prefix “Given a $\psi \in L$ such that $\perp \prec \psi$” is missing from (E4). Secondly the standard definition is actually given relative to an underlying belief set K since it includes, in addition to (E1)–(E4), the following extra condition.

If K is consistent then $\theta \notin K$ iff $\theta \preceq \phi$ for all $\phi \in L$ (minimality)

However, as is noted in [Nay 94], E-relations contain enough information by themselves for this underlying belief set to be extracted from it. Hence we can dispense with (minimality) and instead just define the belief set $Bel(\preceq)$ associated with the E-relation $\preceq$ as follows:

$$Bel(\preceq) = \begin{cases} \{\theta \mid \perp \prec \theta\} & \text{if } \perp \prec \theta, \text{ for some } \theta, \\ L & \text{otherwise.} \end{cases}$$

Hence if $\preceq$ is not absurd, $Bel(\preceq)$ contains precisely those sentences which are strictly more entrenched than $\perp$, while the belief set associated with the absurd E-relation is defined to be the entire set of sentences L. The belief set associated with an E-relation was called its *epistemic content* in [Nay 94].

2.1 E-relations and rational consequence

We now bring in the connection between E-relations, as they have been defined here, and rational consequence relations. The following result is virtually the same as one given in [GM 94]. For this reason we omit the proof.

Proposition 2 *Let $\mid\!\sim$ be a rational consequence relation which is either consistency preserving or trivial. If we define, from $\mid\!\sim$, a binary relation $\preceq_\sim$ on L by setting, for all $\theta, \phi \in L$,*

$$\theta \preceq_\sim \phi \textit{ iff } \neg\theta \vee \neg\phi \not\mid\!\sim \theta \textit{ or } \neg\phi \mid\!\sim \bot, \tag{1}$$

then $\preceq_\sim$ forms an E-relation. Conversely if, given an E-relation $\preceq$ we define a binary relation $\mid\!\sim_\preceq$ on L by setting, for all $\theta, \phi \in L$,

$$\theta \mid\!\sim_\preceq \phi \textit{ iff } \neg\theta \prec \neg\theta \vee \phi \textit{ or } \top \preceq \neg\theta$$

then $\mid\!\sim_\preceq$ forms a rational consequence relation which is either consistency preserving or trivial. Furthermore the identity $\mid\!\sim = \mid\!\sim_{\preceq_\sim}$ holds.

So there is a bijection between rational consequence relations which are either consistency preserving or trivial, and E-relations. Essentially they are different ways of describing the same thing, and so, for example, an operation for changing one automatically gives us an operation for changing the other. This observation is at the heart of the present paper. Given $\vec{\mathcal{U}} \in \Upsilon$ we shall denote by $\preceq_{\vec{\mathcal{U}}}$ the E-relation defined from $\mid\!\sim_{\vec{\mathcal{U}}}$ via (1) above. Since we have already seen that rational consequence relations which are either consistency preserving or trivial are characterised by the sequences in Υ, Proposition 2 immediately leads us to the following result.

Proposition 3 *Let $\preceq$ be a binary relation on L. Then $\preceq$ is an E-relation iff $\preceq = \preceq_{\vec{\mathcal{U}}}$ for some $\vec{\mathcal{U}} \in \Upsilon$.*

Note again that $\preceq_{\vec{\mathcal{U}}} = \preceq_{\vec{\mathcal{V}}}$ does not imply $\vec{\mathcal{U}} = \vec{\mathcal{V}}$. Also note that $\preceq_{\vec{\mathcal{U}}}$ will be absurd iff $\vec{\mathcal{U}}$ is empty. For the case when $\vec{\mathcal{U}}$ is full, we will find it useful to have the following description of $\preceq_{\vec{\mathcal{U}}}$, given directly in terms of $\vec{\mathcal{U}}$. This description can be found by using the definition of $\mid\!\sim_{\vec{\mathcal{U}}}$ via (1).

$$\begin{array}{lll} \theta \preceq_{\vec{\mathcal{U}}} \phi \quad \text{iff} & \textit{either} & \models \phi \\ & \textit{or} & \mathcal{U}_i \cap S_{\neg\theta} \neq \emptyset \text{ for the least } i \text{ such that } \mathcal{U}_i \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset. \end{array}$$

Hence, in terms of plausibility of worlds, $\theta \preceq_{\vec{\mathcal{U}}} \phi$ iff either ϕ is a tautology or, amongst the most plausible worlds (according to $\vec{\mathcal{U}}$) which satisfy either $\neg\theta$ or $\neg\phi$, there is at least one world which satisfies $\neg\theta$. Equivalently the most plausible worlds which satisfy $\neg\theta$ are at least as plausible as the most plausible worlds which satisfy $\neg\phi$. We have the following proposition.

Proposition 4 *Let $\vec{\mathcal{U}} \in \Upsilon$ and $\theta \in L$. Then $\theta \in Bel(\preceq_{\vec{\mathcal{U}}})$ iff $\top \mid\!\sim_{\vec{\mathcal{U}}} \theta$.*

Proof: First suppose $\vec{\mathcal{U}}$ is empty. Then $\preceq_{\vec{\mathcal{U}}}$ is the absurd E-relation and $\mid\!\sim_{\vec{\mathcal{U}}}$ is the trivial rational consequence relation. Hence, in this case, we have that, for *all* θ, both $\theta \in Bel(\preceq_{\vec{\mathcal{U}}})$ and $\top \mid\!\sim_{\vec{\mathcal{U}}} \theta$ and so the result is true. So suppose instead that $\vec{\mathcal{U}}$ is full. Then $\preceq_{\vec{\mathcal{U}}}$ is not absurd. In this case we have $\theta \in Bel(\preceq_{\vec{\mathcal{U}}})$ iff $\perp \prec_{\vec{\mathcal{U}}} \theta$ iff not($\theta \preceq_{\vec{\mathcal{U}}} \perp$). Using the formulation of $\preceq_{\vec{\mathcal{U}}}$ in terms of $\vec{\mathcal{U}}$, this gives us $\theta \in Bel(\preceq_{\vec{\mathcal{U}}})$ iff $\mathcal{U}_i \cap S_{\neg\theta} = \emptyset$ for the least i such that $\mathcal{U}_i \neq \emptyset$. This is easily seen to be equivalent to $\top \mid\!\sim_{\vec{\mathcal{U}}} \theta$ and so the result is proved. ■

Proposition 4 tells us that a sentence is believed in the belief set associated with the E-relation $\preceq_{\vec{\mathcal{U}}}$ iff it is true in all the most plausible worlds according to $\vec{\mathcal{U}}$.

3 Revision of E-relations

In [Nay 94] Nayak deviates from the basic AGM framework in two ways. Firstly, in order to help us deal with *iterated* revision (see [Bou 96, DP 97, Wil 94]), he argues that we need not only a description of the new belief set which results from a revision, but also a new E-relation which can then guide any further revision. Thus we should enlarge our epistemic state to consist of a belief set **together with** an E-relation and then perform revision on this larger state. In fact, since, as we have seen, the belief set may be determined from the E-relation, we may take our epistemic states to be just E-relations.[6] In view of this, for the rest of this paper, we will use the terms "E-relation" and "epistemic state" interchangeably. Secondly, he suggests that the epistemic input should consist not of a single sentence, but rather another E-relation. (See [Nay 94] for motivation.) He claims it is then possible, in his framework, to capture the revision of E-relations by arbitrary sets of sentences E by first converting the set E into a suitable E-relation $\preceq_E$ and then revising by $\preceq_E$. We shall discuss this point further in the next section. In this section we shall use the characterisation of E-relations given in Proposition 3 to describe Nayak's proposal of how one E-relation should be revised by another to obtain a new E-relation. The ideas behind this formulation were also discussed, informally, in [Nay 94] although, for the formal development, Nayak chooses a different formulation.

Let $\preceq_K$ be the prior E-relation and let $\preceq_E$ be the input, or evidential, E-relation. By Proposition 3, we know that there exist $\vec{\mathcal{U}}, \vec{\mathcal{V}} \in \Upsilon$ such that $\preceq_K = \preceq_{\vec{\mathcal{U}}}$ and $\preceq_E = \preceq_{\vec{\mathcal{V}}}$. Hence we may reduce the question of entrenchment revision to a question of how to revise one sequence of world-sets by another. More precisely, we can define a *sequence* revision function $* : \Upsilon \times \Upsilon \to \Upsilon$, where $\vec{\mathcal{U}} * \vec{\mathcal{V}}$ is the result of revising $\vec{\mathcal{U}}$ by $\vec{\mathcal{V}}$, and then simply lift this to an

[6] In this context of iterated revision, the consideration of more comprehensive epistemic states of which a belief set is but one component has also been suggested in [DP 97] and [FH 99].

entrenchment revision function by setting

$$\preceq_K * \preceq_E = \preceq_{\vec{\mathcal{U}}*\vec{\mathcal{V}}} . \tag{2}$$

(The context will always make it clear whether we are considering $*$ as an operation on sequences or an operation on E-relations.) All this, of course, must be independent of precisely which $\vec{\mathcal{U}}$ and $\vec{\mathcal{V}}$ are chosen to represent $\preceq_K$ and $\preceq_E$ respectively. The definition for the sequence revision function $*$ we choose, motivated purely in order to arrive at Nayak's entrenchment revision function, is the following:

Definition 3 *We define the function* $* : \Upsilon \times \Upsilon \to \Upsilon$ *by setting , for all* $\vec{\mathcal{U}} = (\mathcal{U}_0, \ldots, \mathcal{U}_k)$ *and* $\vec{\mathcal{V}} = (\mathcal{V}_0, \ldots, \mathcal{V}_m)$*,*

$$\vec{\mathcal{U}} * \vec{\mathcal{V}} = \begin{cases} (\mathcal{U}_0 \cap \mathcal{V}_0, \mathcal{U}_1 \cap \mathcal{V}_0, \ldots, \mathcal{U}_k \cap \mathcal{V}_0, & \\ \mathcal{U}_0 \cap \mathcal{V}_1, \mathcal{U}_1 \cap \mathcal{V}_1, \ldots, \mathcal{U}_k \cap \mathcal{V}_1, & \text{if } \vec{\mathcal{U}} \text{ is full} \\ \ldots, & \\ \mathcal{U}_0 \cap \mathcal{V}_m, \mathcal{U}_1 \cap \mathcal{V}_m, \ldots, \mathcal{U}_k \cap \mathcal{V}_m). & \\ \vec{\mathcal{V}} & \text{otherwise.} \end{cases}$$

Clearly it is the case that $\vec{\mathcal{U}} * \vec{\mathcal{V}}$ is always full, unless $\vec{\mathcal{V}}$ is empty, in which case so is $\vec{\mathcal{U}} * \vec{\mathcal{V}}$. Hence we certainly have $\vec{\mathcal{U}} * \vec{\mathcal{V}} \in \Upsilon$. Note that the above definition gives us $\vec{\mathcal{U}} * \vec{\mathcal{V}} = (W) * \vec{\mathcal{V}}$ whenever $\vec{\mathcal{U}}$ is empty.

Given $\theta, \phi \in L$, how do we determine whether $\theta \preceq_{\vec{\mathcal{U}}*\vec{\mathcal{V}}} \phi$ holds? If either $\vec{\mathcal{U}}$ or $\vec{\mathcal{V}}$ is empty then we simply have $\preceq_{\vec{\mathcal{U}}*\vec{\mathcal{V}}} = \preceq_{\vec{\mathcal{V}}}$ (either because $\vec{\mathcal{U}} * \vec{\mathcal{V}} = \vec{\mathcal{V}}$ or both $\vec{\mathcal{U}} * \vec{\mathcal{V}}$ and $\vec{\mathcal{V}}$ are empty). But what if both $\vec{\mathcal{U}}$ and $\vec{\mathcal{V}}$ are full? Well if $\models \phi$ then we certainly have $\theta \preceq_{\vec{\mathcal{U}}*\vec{\mathcal{V}}} \phi$, otherwise we must first determine the most plausible worlds, according to $\vec{\mathcal{U}} * \vec{\mathcal{V}}$, in which $\neg\theta \vee \neg\phi$ holds, and then conclude $\theta \preceq_{\vec{\mathcal{U}}*\vec{\mathcal{V}}} \phi$ holds iff at least one of these worlds satisfies $\neg\theta$. Looking at the sequence $\vec{\mathcal{U}} * \vec{\mathcal{V}}$, it is clear that the most plausible worlds which satisfy $\neg\theta \vee \neg\phi$ are precisely those worlds in $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta\vee\neg\phi}$, where j is least such that $\mathcal{V}_j \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset$ and i is then minimal such that $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset$. Thus we have

$$\begin{array}{llll} \theta \preceq_{\vec{\mathcal{U}}*\vec{\mathcal{V}}} \phi & \text{iff} & \textit{either} & \models \phi \\ & & \textit{or} & \mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta} \neq \emptyset \text{ where} \\ & & & \textit{(i).}\ j \text{ is least such that } \mathcal{V}_j \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset \text{ and,} \\ & & & \textit{(ii).}\ i \text{ is then least such that } \mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset \end{array}$$

We would now like to assure ourselves that $*$, when lifted to an operation on E-relations, is well-defined. To do this we make use of the following lemma.

Lemma 2 *Let* $\vec{\mathcal{U}} = (\mathcal{U}_0, \ldots, \mathcal{U}_k)$ *and* $\vec{\mathcal{V}} = (\mathcal{V}_0, \ldots, \mathcal{V}_m)$ *be sequences in* Υ *such that* $\preceq_{\vec{\mathcal{U}}} = \preceq_{\vec{\mathcal{V}}}$*. Given* $\theta \in L$ *let* i *be least such that* $\mathcal{U}_i \cap S_\theta \neq \emptyset$ *and let* j *be least such that* $\mathcal{V}_j \cap S_\theta \neq \emptyset$*. Then* $\mathcal{U}_i \cap S_\theta = \mathcal{V}_j \cap S_\theta$*.*

Proof: Note that it is implicit in the statement of this lemma that $\vec{\mathcal{U}}$ and $\vec{\mathcal{V}}$ are full. Suppose for contradiction that $\mathcal{U}_i \cap S_\theta \neq \mathcal{V}_j \cap S_\theta$, say $w \in \mathcal{U}_i \cap S_\theta$ but $w \notin \mathcal{V}_j \cap S_\theta$ for some $w \in W$. Then it is easy to see that $\neg w \preceq_{\vec{\mathcal{U}}} \neg\theta$ and so, since $\preceq_{\vec{\mathcal{U}}} = \preceq_{\vec{\mathcal{V}}}$, we should also have $\neg w \preceq_{\vec{\mathcal{V}}} \neg\theta$. Since $\vec{\mathcal{V}}$ is full, we know that $w \in \mathcal{V}_{j'}$ for some $0 \leq j' \leq m$ such that $j' \neq j$. But if $j' < j$ then $\mathcal{V}_{j'} \cap S_\theta \neq \emptyset$ – contradicting the minimality of j, while if $j < j'$ then $\neg\theta \prec_{\vec{\mathcal{V}}} \neg w$. Either way we have our required contradiction. ∎

Proposition 5 *Let $\vec{\mathcal{U}}_s, \vec{\mathcal{V}}_s \in \Upsilon$ for $s = 1, 2$. Then $\preceq_{\vec{\mathcal{U}}_1} = \preceq_{\vec{\mathcal{U}}_2}$ and $\preceq_{\vec{\mathcal{V}}_1} = \preceq_{\vec{\mathcal{V}}_2}$ implies $\preceq_{\vec{\mathcal{U}}_1 * \vec{\mathcal{V}}_1} = \preceq_{\vec{\mathcal{U}}_2 * \vec{\mathcal{V}}_2}$.*

Proof: First note that $\preceq_{\vec{\mathcal{U}}_1} = \preceq_{\vec{\mathcal{U}}_2}$ implies that either $\vec{\mathcal{U}}_1$ and $\vec{\mathcal{U}}_2$ are both full or they are both empty. If they are both empty then $\preceq_{\vec{\mathcal{U}}_1 * \vec{\mathcal{V}}_1} = \preceq_{\vec{\mathcal{V}}_1} = \preceq_{\vec{\mathcal{V}}_2} = \preceq_{\vec{\mathcal{U}}_2 * \vec{\mathcal{V}}_2}$ as required. Similarly for $\vec{\mathcal{V}}_1$ and $\vec{\mathcal{V}}_2$. Hence let us assume that all four sequences are full. For each $s = 1, 2$ let us suppose $\vec{\mathcal{U}}_s = (\mathcal{U}^s_0, \ldots, \mathcal{U}^s_{k_s})$ and $\vec{\mathcal{V}} = (\mathcal{V}^s_0, \ldots, \mathcal{V}^s_{m_s})$. Then we have, for each $s = 1, 2$,

$$\begin{array}{lll} \theta \preceq_{\vec{\mathcal{U}}_l * \vec{\mathcal{V}}_l} \phi & \text{iff} & \textit{either} \ \models \phi \\ & & \textit{or} \quad \mathcal{U}^s_{i_s} \cap \mathcal{V}^s_{j_s} \cap S_{\neg\theta} \neq \emptyset \text{ where} \\ & & \qquad \textit{(i).}\ j_s \text{ is least such that } \mathcal{V}^s_{j_s} \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset \text{ and,} \\ & & \qquad \textit{(ii).}\ i_s \text{ is then least such that } \mathcal{U}^s_{i_s} \cap \mathcal{V}^s_{j_s} \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset. \end{array}$$

Then, by Lemma 2 using $\preceq_{\vec{\mathcal{V}}_1} = \preceq_{\vec{\mathcal{V}}_2}$, we have $\mathcal{V}^1_{j_1} \cap S_{\neg\theta\vee\neg\phi} = \mathcal{V}^2_{j_2} \cap S_{\neg\theta\vee\neg\phi}$. Let us momentarily call this set $\mathcal{V}$. Then, since $\mathcal{V} = S_{\bigvee \mathcal{V}}$, we have that, for $s = 1, 2$, i_s is least such that $\mathcal{U}^s_{i_s} \cap S_{\bigvee \mathcal{V}} \neq \emptyset$. Again, by Lemma 2 using $\preceq_{\vec{\mathcal{U}}_1} = \preceq_{\vec{\mathcal{U}}_2}$, we have $\mathcal{U}^1_{i_1} \cap S_{\bigvee \mathcal{V}} = \mathcal{U}^2_{i_2} \cap S_{\bigvee \mathcal{V}}$. Thus, wrapping all this together and noting that $S_{\neg\theta} = S_{\neg\theta\vee\neg\phi} \cap S_{\neg\theta}$ we have $\mathcal{U}^1_{i_1} \cap \mathcal{V}^1_{j_1} \cap S_{\neg\theta} = (\mathcal{U}^1_{i_1} \cap \mathcal{V}^1_{j_1} \cap S_{\neg\theta\vee\neg\phi}) \cap S_{\neg\theta} = (\mathcal{U}^1_{i_1} \cap S_{\bigvee \mathcal{V}}) \cap S_{\neg\theta} = (\mathcal{U}^2_{i_2} \cap S_{\bigvee \mathcal{V}}) \cap S_{\neg\theta} = (\mathcal{U}^2_{i_2} \cap \mathcal{V}^2_{j_2} \cap S_{\neg\theta\vee\neg\phi}) \cap S_{\neg\theta} = \mathcal{U}^2_{i_2} \cap \mathcal{V}^2_{j_2} \cap S_{\neg\theta}$. The result follows. ∎

Thus, by Proposition 5, the operation $\preceq_K * \preceq_E$ is indeed independent of which $\vec{\mathcal{U}}, \vec{\mathcal{V}}$ we choose such that $\preceq_K = \preceq_{\vec{\mathcal{U}}}$ and $\preceq_E = \preceq_{\vec{\mathcal{V}}}$.

From now on we will follow Nayak and use $\preceq_{K*E}$ as an abbreviation for $\preceq_K * \preceq_E$. In [NNP 96] Nayak et al propose the following postulates for the revision of E-relations:[7]

[7] Actually, the list of postulates given in [NNP 96] differs from the list here in that the prefix "For $\preceq_E$ non-absurd" is missing from (E4*), while (E2*) is missing completely. It seems that this difference is due to nothing more than a small oversight by the authors of [NNP 96] and that this list corresponds to what they actually intended. Indeed this has been confirmed by Nayak in personal communication. Since the damage caused by this oversight is essentially confined only to the special limiting case when the input E-relation is absurd, it does not have any serious consequences for the results presented in [NNP 96]. In particular only slight modifications to their proof are necessary to show that Nayak's operation of revision is characterised by the list of postulates given here. Hence to avoid complicating our discussion any further, we will simply carry on *as though* this is the list given in [NNP 96].

(E1*) $\preceq_{K*E}$ is an E-relation.

(E2*) If $\preceq_E$ is absurd then so is $\preceq_{K*E}$.

(E3*) If $\theta \prec_E \phi$ then $\theta \prec_{K*E} \phi$.

(E4*) For $\preceq_E$ non-absurd, if both $\theta \preceq_E \phi$ and $\phi \preceq_E \theta$ and if, for all λ, χ such that $\theta \wedge \phi \models \chi$ and $\theta \prec_E \chi$, we have $\lambda \preceq_K \chi$ iff $\lambda \preceq_E \chi$, then $\theta \preceq_{K*E} \phi$ iff $\theta \preceq_K \phi$.

Any operation of revision of E-relations which satisfies the above four conditions is called a *well-behaved* entrenchment revision operation in [NNP 96]. (E1*) is a minimal requirement for $\preceq_{K*E}$, while (E2*) takes care of the limiting case when the input E-relation is absurd. (E3*) expresses the intuition that, upon revision, evidence should take *priority*. That is, whenever the evidence says that ϕ is strictly more entrenched than θ, then so too must the revised E-relation. In terms of sequences of world-sets, (E3*) says, roughly, that the sequence $\vec{\mathcal{U}} * \vec{\mathcal{V}}$ associated with the revised E-relation should be a "refinement" or "thinning" of the evidential sequence $\vec{\mathcal{V}}$, i.e., we sub-partition each $\mathcal{V}_j$ into some $\vec{\mathcal{W}}^j = (\mathcal{W}_0^j, \ldots, \mathcal{W}_{m_j}^j)$ and then (abusing notation slightly) let $\vec{\mathcal{U}} * \vec{\mathcal{V}} = (\vec{\mathcal{W}}^0, \ldots, \vec{\mathcal{W}}^m)$. For a justification of (E4*) we refer the reader to [NNP 96]. This postulate looks complicated, but its effect is merely to tell us exactly how the sub-partition mentioned above should take place. We say exactly, since it is shown in [NNP 96] that there is precisely one revision operation on E-relations which satisfies (E1*)–(E4*), namely the one given in [Nay 94]. Thus the above four postulates serve to characterise Nayak's method. We show that our revision operation, defined by Definition 3 via (2) above, also satisfies (E1*)–(E4*), thus confirming its equivalence with the operation constructed in [Nay 94].

Theorem 1 *If we set $\preceq_{K*E} = \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}}$ where $\vec{\mathcal{U}}$ ($\vec{\mathcal{V}}$) is chosen so that $\preceq_K = \preceq_{\vec{\mathcal{U}}}$ ($\preceq_E = \preceq_{\vec{\mathcal{V}}}$) then the operator $*$ satisfies (E1*), (E2*), (E3*) and (E4*).*

Proof: Let $\vec{\mathcal{U}} = (\mathcal{U}_0, \ldots, \mathcal{U}_k)$ and $\vec{\mathcal{V}} = (\mathcal{V}_0, \ldots, \mathcal{V}_m)$ be such that $\preceq_K = \preceq_{\vec{\mathcal{U}}}$ and $\preceq_E = \preceq_{\vec{\mathcal{V}}}$. That $\preceq_{K*E}$ satisfies (E1*) is clear. To show (E2*) suppose $\preceq_E$ is absurd. Then $\vec{\mathcal{V}}$ is empty and hence so is $\vec{\mathcal{U}} * \vec{\mathcal{V}}$. Thus $\preceq_{K*E} = \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}}$ is also absurd as required. To show (E3*) we must show that $\theta \prec_{\vec{\mathcal{V}}} \phi$ implies $\theta \prec_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \phi$, equivalently $\phi \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \theta$ implies $\phi \preceq_{\vec{\mathcal{V}}} \theta$. So suppose $\phi \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \theta$. If either $\vec{\mathcal{U}}$ or $\vec{\mathcal{V}}$ are empty then we have $\preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} = \prec_{\vec{\mathcal{V}}}$ and so $\phi \preceq_{\vec{\mathcal{V}}} \theta$ as required. So assume both $\vec{\mathcal{U}}$ and $\vec{\mathcal{V}}$ are full. If $\models \theta$ then again $\phi \preceq_{\vec{\mathcal{V}}} \theta$ as required, so suppose also $\not\models \theta$. Let j be least such that $\mathcal{V}_j \cap S_{\neg\phi \vee \neg\theta} \neq \emptyset$ and let i be least such that $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\phi \vee \neg\theta} \neq \emptyset$. Then we have $\phi \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \theta$ implies $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\phi} \neq \emptyset$. Hence $\mathcal{V}_j \cap S_{\neg\phi} \neq \emptyset$ and so $\phi \preceq_{\vec{\mathcal{V}}} \theta$, thus proving (E3*).
To show (E4*), first note that, by Lemma 1, we may replace the occurrence of "$\theta \prec_E \chi$" in this postulate by "$\theta \wedge \phi \prec_E \chi$". Hence showing $\preceq_{K*E}$ satisfies (E4*) is equivalent to showing that $\preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}}$ satisfies:

(A^*) For $\preceq_{\vec{\mathcal{V}}}$ non-absurd, if both $\theta \preceq_{\vec{\mathcal{V}}} \phi$ and $\phi \preceq_{\vec{\mathcal{V}}} \theta$ and if, for all λ, χ such that $\theta \wedge \phi \models \chi$ and $\theta \wedge \phi \prec_{\vec{\mathcal{V}}} \chi$, we have $\lambda \preceq_{\vec{\mathcal{U}}} \chi$ iff $\lambda \preceq_{\vec{\mathcal{V}}} \chi$, then $\theta \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \phi$ iff $\theta \preceq_{\vec{\mathcal{U}}} \phi$.

If $\vec{\mathcal{V}}$ is empty then $\preceq_{\vec{\mathcal{V}}}$ is absurd and so (A^*) holds vacuously, so assume $\vec{\mathcal{V}}$ is full. Then if $\vec{\mathcal{U}}$ is empty we have $\vec{\mathcal{U}}$ is absurd and $\vec{\mathcal{U}} * \vec{\mathcal{V}} = \vec{\mathcal{V}}$, so showing the conclusion "$\theta \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \phi$ iff $\theta \preceq_{\vec{\mathcal{U}}} \phi$" of (A^*) reduces to showing "$\theta \preceq_{\vec{\mathcal{V}}} \phi$". But this condition already appears in the hypotheses of (A^*) and so (A^*) also holds in this case. So now let us suppose that both $\vec{\mathcal{U}}$ and $\vec{\mathcal{V}}$ are full. Let θ and ϕ satisfy the hypotheses of (A^*). If $\models \phi$ then we automatically have both $\theta \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \phi$ and $\theta \preceq_{\vec{\mathcal{U}}} \phi$ as required. So let us assume that $\not\models \phi$. Let j be minimal such that $\mathcal{V}_j \cap S_{\neg\theta \vee \neg\phi} \neq \emptyset$. Note that, since $\theta \preceq_{\vec{\mathcal{V}}} \phi$, we have $\mathcal{V}_j \cap S_{\neg\theta} \neq \emptyset$. Let i be minimal such that $\mathcal{U}_i \cap S_{\neg\theta \vee \neg\phi} \neq \emptyset$. We claim that $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta \vee \neg\phi} \neq \emptyset$ and thus that i is also minimal such that this holds. To prove this suppose for contradiction that $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta \vee \neg\phi} = \emptyset$. Then, as can easily be checked, this would give us

$$\theta \wedge \phi \prec_{\vec{\mathcal{V}}} (\theta \wedge \phi) \vee \neg \bigvee \mathcal{U}_i.$$

We also have

$$(\theta \wedge \phi) \vee \neg \bigvee \mathcal{U}_i \prec_{\vec{\mathcal{U}}} (\theta \wedge \phi) \vee \neg \bigvee \mathcal{V}_j$$

and

$$(\theta \wedge \phi) \vee \neg \bigvee \mathcal{V}_j \preceq_{\vec{\mathcal{V}}} (\theta \wedge \phi) \vee \neg \bigvee \mathcal{U}_i.$$

Hence, putting $\lambda = (\theta \wedge \phi) \vee \neg \bigvee \mathcal{V}_j$ and $\chi = (\theta \wedge \phi) \vee \neg \bigvee \mathcal{U}_i$, we have shown the existence of sentences λ, χ such that $\theta \wedge \phi \models \chi$, $\theta \wedge \phi \prec_{\vec{\mathcal{V}}} \chi$, $\lambda \preceq_{\vec{\mathcal{V}}} \chi$ but $\chi \prec_{\vec{\mathcal{U}}} \lambda$. This contradicts the hypotheses of (A^*) and so we must have $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta \vee \neg\phi} \neq \emptyset$ as required. Now to show that $\theta \preceq_{\vec{\mathcal{U}} * \vec{\mathcal{V}}} \phi$ iff $\theta \preceq_{\vec{\mathcal{U}}} \phi$ it remains to show that $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta} = \emptyset$ iff $\mathcal{U}_i \cap S_{\neg\theta} = \emptyset$. The "if" clause here is immediate. To show the "only if" clause, suppose for contradiction that $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta} = \emptyset$ but $\mathcal{U}_i \cap S_{\neg\theta} \neq \emptyset$. Then $\mathcal{U}_i \cap \mathcal{V}_j \cap S_{\neg\theta} = \emptyset$ by itself gives us

$$\theta \wedge \phi \prec_{\vec{\mathcal{V}}} \theta \vee \neg \bigvee \mathcal{U}_i,$$

while, together with $\mathcal{U}_i \cap S_{\neg\theta} \neq \emptyset$, it gives us

$$\theta \vee \neg \bigvee \mathcal{U}_i \prec_{\vec{\mathcal{U}}} \theta \vee \neg \bigvee \mathcal{V}_j.$$

Meanwhile, using the fact that $\mathcal{V}_j \cap S_{\neg\theta} \neq \emptyset$ (since $\theta \preceq_{\vec{\mathcal{V}}} \phi$), we get

$$\theta \vee \neg \bigvee \mathcal{V}_j \preceq_{\vec{\mathcal{V}}} \theta \vee \neg \bigvee \mathcal{U}_i.$$

Hence, this time putting $\lambda = \theta \vee \neg \bigvee \mathcal{V}_j$ and $\chi = \theta \vee \neg \bigvee \mathcal{U}_i$, we have again found sentences λ and χ such that $\theta \wedge \phi \models \chi, \theta \wedge \phi \prec_{\vec{\mathcal{V}}} \chi, \lambda \preceq_{\vec{\mathcal{V}}} \chi$ but $\chi \prec_{\vec{\mathcal{U}}} \lambda$. This contradicts the hypotheses of (A^*) and so we must have $\mathcal{U}_i \cap S_{\neg\theta} = \emptyset$ as required. This completes the proof. ■

One advantage of this particular formulation is that it is relatively easy to show properties of the well-behaved entrenchment revision operation $*$. For example, the following proposition regarding sequence revision is straightforward to prove.

Proposition 6 *Let $\vec{\mathcal{U}}, \vec{\mathcal{V}}, \vec{\mathcal{W}} \in \Upsilon$ and suppose $\vec{\mathcal{V}}$ is not empty. Then $(\vec{\mathcal{U}} * \vec{\mathcal{V}}) * \vec{\mathcal{W}} = \vec{\mathcal{U}} * (\vec{\mathcal{V}} * \vec{\mathcal{W}})$.*

Proof: First of all suppose $\vec{\mathcal{U}}$ is empty. Then $\vec{\mathcal{U}} * \vec{\mathcal{V}} = \vec{\mathcal{V}}$ and so $(\vec{\mathcal{U}} * \vec{\mathcal{V}}) * \vec{\mathcal{W}} = \vec{\mathcal{V}} * \vec{\mathcal{W}} = \vec{\mathcal{U}} * (\vec{\mathcal{V}} * \vec{\mathcal{W}})$ as required. So suppose $\vec{\mathcal{U}}$ is not empty. Then, using the assumption that $\vec{\mathcal{V}}$ is not empty, neither is $\vec{\mathcal{U}} * \vec{\mathcal{V}}$. In this case, assuming $\vec{\mathcal{U}} = (\mathcal{U}_0, \ldots, \mathcal{U}_k), \vec{\mathcal{V}} = (\mathcal{V}_0, \ldots, \mathcal{V}_m)$ and $\vec{\mathcal{W}} = (\mathcal{W}_0, \ldots, \mathcal{W}_l)$, it is easy but tedious to check that both $(\vec{\mathcal{U}} * \vec{\mathcal{V}}) * \vec{\mathcal{W}}$ and $\vec{\mathcal{U}} * (\vec{\mathcal{V}} * \vec{\mathcal{W}})$ come out to be the sequence of length $(k+1) \times (m+1) \times (l+1)$ whose elements are all the sets of the form $\mathcal{U}_a \cap \mathcal{V}_b \cap \mathcal{W}_c$ where $0 \leq a \leq k$, $0 \leq b \leq m$ and $0 \leq c \leq l$, and in which $\mathcal{U}_{a_1} \cap \mathcal{V}_{b_1} \cap \mathcal{W}_{c_1}$ will appear before $\mathcal{U}_{a_2} \cap \mathcal{V}_{b_2} \cap \mathcal{W}_{c_2}$ iff either $c_1 < c_2$, or $c_1 = c_2$ and $b_1 < b_2$, or $c_1 = c_2, b_1 = b_2$ and $a_1 < a_2$. ■

Note that, in general, we do need the condition on $\vec{\mathcal{V}}$ here, since if $\vec{\mathcal{V}}$ is empty then so is $(\vec{\mathcal{U}} * \vec{\mathcal{V}})$ and so $(\vec{\mathcal{U}} * \vec{\mathcal{V}}) * \vec{\mathcal{W}} = \vec{\mathcal{W}}$, while $\vec{\mathcal{U}} * (\vec{\mathcal{V}} * \vec{\mathcal{W}}) = \vec{\mathcal{U}} * \vec{\mathcal{W}}$. This proposition, in turn, immediately gives us the following interesting associativity property of the induced entrenchment revision operation.

Proposition 7 *Let $\preceq_i$ be an E-relation for $i = 1, 2, 3$. Then, if $\preceq_2$ is not absurd, we have $(\preceq_1 * \preceq_2) * \preceq_3 = \preceq_1 * (\preceq_2 * \preceq_3)$.*

We now turn to the question of how to generate an E-relation from a set of sentences.

4 Generating E-relations from sets of sentences

As we said in the last section, Nayak proposes that his way of revising one E-relation by another allows a way of modelling the revision of an E-relation by a set of sentences E by first converting, according to some suitable method, the set E into an E-relation $\preceq_E$ and then revising by $\preceq_E$. The question of which "suitable method" we should use for generating $\preceq_E$ is clearly an interesting question in itself. A strong feeling is that the relation $\preceq_E$ should adequately convey the informational content of E, but what does this mean? An obvious first requirement of $\preceq_E$ would seem to be $Bel(\preceq_E) = Cn(E)$, but there are different ways in which this can be achieved. The definition which Nayak seems to advocate is the following, based on an idea of Rott [Rot 92a], and expressed via its strict part.

$\theta \prec_E \phi$ iff $E \not\models \bot, \not\models \theta$ and for all $E' \subseteq E$ such that $E' \cup \{\neg\phi\}$ is consistent, there exists $E'' \subseteq E$ such that $E' \subset E''$ and $E'' \cup \{\neg\theta\}$ is consistent.

The clause "$E \not\models \bot$" in the above merely ensures that if E is inconsistent then $\preceq_E$ is absurd, while the clause "$\not\models \theta$" ensures that tautologies are maximally entrenched. The main body of the definition says that ϕ should be strictly more entrenched than θ iff every subset of E which is consistent with $\neg\phi$ can be strictly enlarged to a subset of E which is consistent with $\neg\theta$, or, to put it another way, each $\subseteq$-maximal subset of E which fails to imply ϕ may be strictly enlarged to a subset of E which fails to imply θ. The problem with defining $\preceq_E$ in this way is that it will fail, in general, to be an E-relation. In particular it will not necessarily satisfy (E1).[8] How can we modify/extend $\preceq_E$ above so as to obtain an E-relation? The possibility we choose is to compare the sets which fail to imply θ and ϕ by cardinality rather than inclusion:[9]

Definition 4 *Given a set $E \subseteq L$, define a relation $\prec_E \subseteq L \times L$ by, for all $\theta, \phi \in L$,*

$\theta \prec_E \phi$ *iff* $E \not\models \bot, \not\models \theta$ *and for all $E' \subseteq E$ such that $E' \cup \{\neg\phi\}$ is consistent, there exists $E'' \subseteq E$ such that $|E'| < |E''|$ and $E'' \cup \{\neg\theta\}$ is consistent.*

Note that this definition does indeed extend the "old" definition given above. That $\preceq_E$ defined by Definition 4 is a genuine E-relation will follow once we have found a sequence $\vec{\mathcal{U}} \in \Upsilon$ such that $\preceq_E = \preceq_{\vec{\mathcal{U}}}$. We do this as follows. Let us assume for simplicity that E is finite with $|E| = k$. Then, for each $i = 0, \ldots, k$, we set

$$\mathcal{U}_i^E = \begin{cases} \{w \in W \mid |sent_E(w)| = k - i\} & \text{if } E \not\models \bot \\ \emptyset & \text{otherwise.} \end{cases}$$

So, in the principal case when E is consistent, $\mathcal{U}_i^E$ contains those worlds which satisfy precisely $k - i$ elements of E. Let $\vec{\mathcal{U}}^E = (\mathcal{U}_0^E, \ldots, \mathcal{U}_k^E)$.

Proposition 8 *If $E \models \bot$ then $\vec{\mathcal{U}}^E$ is empty, while if $E \not\models \bot$ then $\vec{\mathcal{U}}^E$ is full (and so, either way, $\vec{\mathcal{U}}^E \in \Upsilon$). In both cases we have $\preceq_E = \preceq_{\vec{\mathcal{U}}^E}$. Hence $\preceq_E$ is an E-relation.*

Proof: The first part of the proposition is obvious, so let us concentrate on showing $\preceq_E = \preceq_{\vec{\mathcal{U}}^E}$. We will first show that $\theta \prec_{\vec{\mathcal{U}}^E} \phi$ implies $\theta \prec_E \phi$. So suppose $\theta \prec_{\vec{\mathcal{U}}^E} \phi$. Then $\prec_{\vec{\mathcal{U}}^E}$ cannot be absurd and so $\vec{\mathcal{U}}^E$ must be full which, in turn, means $E \not\models \bot$. Clearly also $\not\models \theta$, so it remains to show that for each subset of E which is consistent with $\neg\phi$, there exists a subset of E which contains strictly more elements and which is consistent with $\neg\theta$. Let i be minimal such that $\mathcal{U}_i^E \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset$. Then $\theta \prec_{\vec{\mathcal{U}}^E} \phi$ implies $\mathcal{U}_i^E \cap S_{\neg\phi} = \emptyset$ and so we must have $\mathcal{U}_i^E \cap S_{\neg\theta} \neq \emptyset$, while $\mathcal{U}_j^E \cap S_{\neg\phi} = \emptyset$ for all $j \leq i$. Let $w_0 \in \mathcal{U}_i^E \cap S_{\neg\theta}$. Then

[8] It should be noted, however, that $\preceq_E$ so defined does still enjoy several interesting properties. In fact it belongs to Rott's family of *generalized* E-relations [Rot 92b].

[9] Possibilities in this spirit are also discussed in [BCDLP 93] (see Section 2 on "flat belief bases") and in [Leh 95] (see Section 8 on "competing but equal defaults"). See also the closely related Section 5 of [Fre 99].

we have that $sent_E(w_0) \cup \{\neg\theta\}$ is consistent, while $|sent_E(w_0)| = k - i$. Now let $E' \subseteq E$ be such that $E' \cup \{\neg\phi\}$ is consistent and let $w'_0 \in W$ be such that $w'_0 \models \bigwedge E' \wedge \neg\phi$. Then clearly $|sent_E(w'_0)| \geq |E'|$ (w'_0 may satisfy more sentences in E than just those in E'). We need to show that there is some $E'' \subseteq E$ such that $E'' \cup \{\neg\theta\}$ is consistent and $|E'| < |E''|$. Try $E'' = sent_E(w_0)$. If we had $|E'| \geq k - i$ then we would have $|sent_E(w'_0)| \geq |E'| \geq k - i$. Thus we would get $\mathcal{U}^E_j \cap S_{\neg\phi} \neq \emptyset$ for some $j \leq i$ (take $j = k - |sent_E(w'_0)|$) – contradiction. Hence $|E'| < k - i = |sent_E(w_0)|$ as required.
For the converse, suppose $\theta \prec_E \phi$. Then $E \not\models \bot$ – which gives us that $\vec{\mathcal{U}}^E$ is full – and $\not\models \theta$. Let i be least such that $\mathcal{U}^E_i \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset$. We must show $\mathcal{U}^E_i \cap S_{\neg\phi} = \emptyset$. But suppose for contradiction that $\mathcal{U}^E_i \cap S_{\neg\phi} \neq \emptyset$ and that $w_0 \in \mathcal{U}^E_i \cap S_{\neg\phi}$ for some $w_0 \in W$. Then $sent_E(w_0) \cup \{\neg\phi\}$ is consistent. Hence, since $\theta \prec_E \phi$, there is some $E'' \subseteq E$ such that $E'' \cup \{\neg\theta\}$ is consistent and $|E''| > |sent_E(w_0)| = k - i$ (since $w_0 \in \mathcal{U}^E_i$). Choose $w'_0 \in W$ such that $w'_0 \models \bigwedge E'' \wedge \neg\theta$. Then $|sent_E(w'_0)| \geq |E''|$ and $w'_0 \in S_{\neg\theta\vee\neg\phi}$. Thus there exists $j < i$ such that $\mathcal{U}^E_j \cap S_{\neg\theta\vee\neg\phi} \neq \emptyset$ (take $j = k - |sent_E(w'_0)|$). This contradicts the minimality of i and so we must have $\mathcal{U}^E_i \cap S_{\neg\phi} = \emptyset$ as required to show $\theta \prec_{\vec{\mathcal{U}}^E} \phi$. ■

Note that, with the notation given above, we have $\vec{\mathcal{U}}^\emptyset = (W)$. Hence we can think of $\preceq_\emptyset$ as being the maximally ignorant epistemic state in which each world is judged to be equally plausible.

How does $\preceq_E$ portray the informational content of E? The sequence $\vec{\mathcal{U}}^E$ shows us clearly. First of all it is easy to see that $\preceq_E$ satisfies the basic requirement of $Bel(\preceq_E) = Cn(E)$ (in particular the only sentences believed in $\preceq_\emptyset$ are the tautologies) since the most plausible worlds in $\vec{\mathcal{U}}^E$, i.e., the worlds in $\mathcal{U}^E_0$, are precisely those worlds which satisfy every sentence in E. The big question is how does $\vec{\mathcal{U}}^E$ classify the worlds which do **not** satisfy every sentence in E. The answer is that it considers one such world more plausible than another iff it satisfies strictly more of the sentences in E. This makes the relation $\preceq_E$ dependent on the syntactic form, not just the semantic form, of E, i.e., **we can have** $Cn(E_1) = Cn(E_2)$ **without necessarily having** $\preceq_{E_1} = \preceq_{E_2}$. For example it is not generally the case that $\preceq_{\{\theta,\phi\}} = \preceq_{\{\theta\wedge\phi\}}$. One situation where this method might be deemed suitable is if we want to regard the elements of E as items of information coming from different, independent sources. A consequence of this is that if we identify $\preceq *E$ with $\preceq * \preceq_E$, then we do not necessarily have $Cn(E_1) = Cn(E_2)$ implies $\preceq *E_1 = \preceq *E_2$. However, as can easily be checked, we *do* have $Cn(E_1) = Cn(E_2)$ implies $Bel(\preceq *E_1) = Bel(\preceq *E_2)$. Hence this sensitivity to the syntactic form of E will reveal itself only in iterated revision.

From now on, for the special case when E is a singleton, we shall write $\preceq_\theta$ rather than $\preceq_{\{\theta\}}$ and $\vec{\mathcal{U}}^\theta$ rather than $\vec{\mathcal{U}}^{\{\theta\}}$. We have the following partial generalisation of Proposition 4.

Proposition 9 *Let $\vec{\mathcal{U}} \in \Upsilon$ be full and let $\theta, \phi \in L$. Then $\phi \in Bel(\preceq_{\vec{\mathcal{U}}} * \preceq_\theta)$ iff $\theta \mid\!\sim_{\vec{\mathcal{U}}} \phi$.*

Proof: From Proposition 8 we have that $\preceq_{\vec{\mathcal{U}}} * \preceq_\theta = \preceq_{\vec{\mathcal{U}}} * \preceq_{\vec{\mathcal{U}}^\theta}$ which, in turn, is equal to $\preceq_{\vec{\mathcal{U}}*\vec{\mathcal{U}}^\theta}$. Suppose first of all that θ is inconsistent. Then $\vec{\mathcal{U}}^\theta$ is empty and hence so is $\vec{\mathcal{U}}*\vec{\mathcal{U}}^\theta$. Thus $\preceq_{\vec{\mathcal{U}}*\vec{\mathcal{U}}^\theta}$ is equal to the absurd E-relation and so we have that $\phi \in Bel(\preceq_{\vec{\mathcal{U}}} * \preceq_\theta) = Bel(\preceq_{\vec{\mathcal{U}}*\vec{\mathcal{U}}^\theta})$ for all ϕ. Meanwhile θ inconsistent also implies $\theta \mid\!\sim_{\vec{\mathcal{U}}} \phi$ for all ϕ. Hence the result is true in this case. So suppose now θ is consistent. Then $\vec{\mathcal{U}}^\theta = (S_\theta, S_{\neg\theta})$ and so, assuming $\vec{\mathcal{U}} = (\mathcal{U}_0, \ldots, \mathcal{U}_k)$ we have that $\vec{\mathcal{U}} * \vec{\mathcal{U}}^\theta = (\mathcal{U}_0 \cap S_\theta, \ldots, \mathcal{U}_k \cap S_\theta, \mathcal{U}_0 \cap S_{\neg\theta}, \ldots, \mathcal{U}_k \cap S_{\neg\theta})$. By Proposition 4 we have that $\phi \in Bel(\preceq_{\vec{\mathcal{U}}*\vec{\mathcal{U}}^\theta})$ iff $\top \mid\!\sim_{\vec{\mathcal{U}}*\vec{\mathcal{U}}^\theta} \phi$ which means that $\phi \in Bel(\preceq_{\vec{\mathcal{U}}*\vec{\mathcal{U}}^\theta})$ iff $\mathcal{U}_i \cap S_\theta \subseteq S_\phi$ for the least i such that $\mathcal{U}_i \cap S_\theta \neq \emptyset$. This is easily seen to be equivalent to $\theta \mid\!\sim_{\vec{\mathcal{U}}} \phi$, thus completing the proof. ∎

Note we require $\vec{\mathcal{U}}$ to be full in this proposition. The result does not hold for empty $\vec{\mathcal{U}}$ since, in this case, we have $\theta \mid\!\sim_{\vec{\mathcal{U}}} \phi$ for all θ, ϕ while $\phi \in Bel(\preceq_{\vec{\mathcal{U}}} * \preceq_\theta)$ iff $\theta \models \phi$.

We are now ready to give the sequence $\vec{\mathcal{U}}$ such that $\theta \mid\!\sim_{\vec{\mathcal{U}}} \phi$ iff $\theta \mid\!\sim^\Delta_{lex} \phi$. Let $(\Delta_0, \ldots, \Delta_n)$ be the Z-partition of Δ. Then, to obtain our special $\vec{\mathcal{U}}$ we start at the sequence (W) and then successively revise, using our sequence revision function $*$, by $\vec{\mathcal{U}}^{\Delta_i^\rightarrow}$ for $i = 0, 1, \ldots, n$, i.e., recalling that $(W) = \vec{\mathcal{U}}^\emptyset$, we have $\theta \mid\!\sim^\Delta_{lex} \phi$ iff $\theta \mid\!\sim_{\vec{\mathcal{U}}^\emptyset * \vec{\mathcal{U}}^{\Delta_0^\rightarrow} * \cdots * \vec{\mathcal{U}}^{\Delta_n^\rightarrow}} \phi$. Note that, by Proposition 6 and the assumption that $\Delta^\rightarrow$ is consistent, the term $\vec{\mathcal{U}}^\emptyset * \vec{\mathcal{U}}^{\Delta_0^\rightarrow} * \cdots * \vec{\mathcal{U}}^{\Delta_n^\rightarrow}$ here is independent of the bracketing. The following extra piece of notation will help us to prove this.

Definition 5 *We define a binary relation $\lhd_{lex}$ on the set of $(n+1)$-tuples $\{\langle i_0, \ldots, i_n\rangle \mid i_r \in \{0, 1, \ldots, |\Delta_r|\}$ for $r = 0, \ldots, n\}$ by*

$$\langle i_0, \ldots, i_n\rangle \lhd_{lex} \langle j_0, \ldots, j_n\rangle \quad \textit{iff} \quad \begin{array}{l} \textit{there exists } r \textit{ such that } i_r < j_r \textit{ and,} \\ \textit{for all } s > r,\ i_s = j_s. \end{array}$$

We can use this notation to help us describe the relation $\mid\!\sim_{\vec{\mathcal{U}}^\emptyset * \vec{\mathcal{U}}^{\Delta_0^\rightarrow} * \cdots * \vec{\mathcal{U}}^{\Delta_n^\rightarrow}}$ as follows: It may be checked that $\vec{\mathcal{U}}^\emptyset * \vec{\mathcal{U}}^{\Delta_0^\rightarrow} * \cdots * \vec{\mathcal{U}}^{\Delta_n^\rightarrow} = \vec{\mathcal{U}}^{\Delta_0^\rightarrow} * \cdots * \vec{\mathcal{U}}^{\Delta_n^\rightarrow}$ works out to be the sequence of length $\prod_{i=0}^n (|\Delta_i| + 1)$ whose elements are all the sets of the form $\mathcal{U}_{s_0}^{\Delta_0^\rightarrow} \cap \mathcal{U}_{s_1}^{\Delta_1^\rightarrow} \cap \ldots \cap \mathcal{U}_{s_n}^{\Delta_n^\rightarrow}$ where $0 \leq s_i \leq |\Delta_i|$ for each $i = 0, 1, \ldots, n$, and in which $\mathcal{U}_{s_0}^{\Delta_0^\rightarrow} \cap \ldots \cap \mathcal{U}_{s_n}^{\Delta_n^\rightarrow}$ appears before $\mathcal{U}_{t_0}^{\Delta_0^\rightarrow} \cap \ldots \cap \mathcal{U}_{t_n}^{\Delta_n^\rightarrow}$ iff $\langle s_0, \ldots, s_n\rangle \lhd_{lex} \langle t_0, \ldots, t_n\rangle$. Thus we may write

$$\theta \mid\!\sim_{\vec{\mathcal{U}}^\emptyset * \vec{\mathcal{U}}^{\Delta_0^\rightarrow} * \cdots * \vec{\mathcal{U}}^{\Delta_n^\rightarrow}} \phi \quad \text{iff} \quad \begin{array}{ll} \textit{either} & \theta \models \bot \\ \textit{or} & \text{for } \langle i_0, \ldots, i_n\rangle \text{ minimal under } \lhd_{lex} \text{ such} \\ & \text{that } \mathcal{U}_{i_0}^{\Delta_0^\rightarrow} \cap \ldots \cap \mathcal{U}_{i_n}^{\Delta_n^\rightarrow} \cap S_\theta \neq \emptyset, \text{ we have} \\ & \mathcal{U}_{i_0}^{\Delta_0^\rightarrow} \cap \ldots \cap \mathcal{U}_{i_n}^{\Delta_n^\rightarrow} \cap S_\theta \subseteq S_\phi. \end{array}$$

We may also use the relation $\lhd_{lex}$ to slightly re-phrase our definition of $\mid\!\sim^\Delta_{lex}$. Recall that

$$\theta \mid\!\sim^\Delta_{lex} \phi \quad \text{iff} \quad \begin{array}{l} \text{for all } \Gamma \subseteq \Delta^\rightarrow \text{ such that } \Gamma \cup \{\theta\} \text{ is consistent and } \Gamma \text{ is} \\ \ll_{lex}\text{-maximal amongst such subsets, we have } \Gamma \cup \{\theta\} \models \phi. \end{array}$$

The relation $\ll_{lex}$ can now be written in terms of $\lhd_{lex}$ by, for each $A, B \subseteq \Delta^{\rightarrow}$,

$$A \ll_{lex} B \text{ iff } \langle |A_0|, \ldots, |A_n| \rangle \lhd_{lex} \langle |B_0|, \ldots, |B_n| \rangle,$$

where, for each $r = 0, 1, \ldots, n$, $A_i = A \cap \Delta_i^{\rightarrow}$ and $B_i = B \cap \Delta_i^{\rightarrow}$. The following property of $\lhd_{lex}$ is easily seen to be true.

Lemma 3 *Let $i_r, j_r \in \{0, 1, \ldots, |\Delta_r|\}$ for $r = 0, 1, \ldots, n$. Then*

$$\langle i_0, \ldots, i_n \rangle \lhd_{lex} \langle j_0, \ldots, j_n \rangle \text{ iff } \langle |\Delta_0| - j_0, \ldots, |\Delta_n| - j_n \rangle \lhd_{lex} \langle |\Delta| - i_0, \ldots, |\Delta_n| - i_n \rangle.$$

With this notation in place we may now give our main result.

Theorem 2 *Let Δ be set of defaults with associated Z-partition $(\Delta_0, \ldots, \Delta_n)$. Then, for all $\theta, \phi \in L$, we have $\theta \mid\!\sim^{\Delta}_{lex} \phi$ iff $\theta \mid\!\sim_{\vec{\mathcal{U}}^{\emptyset} * \vec{\mathcal{U}}^{\Delta_0^{\rightarrow}} * \cdots * \vec{\mathcal{U}}^{\Delta_n^{\rightarrow}}} \phi$.*

Proof: Throughout the proof, given any $\Gamma \subseteq \Delta^{\rightarrow}$, we let $\Gamma_i = \Gamma \cap \Delta_i^{\rightarrow}$. Let us show that $\theta \mid\!\sim^{\Delta}_{lex} \phi$ implies $\theta \mid\!\sim_{\vec{\mathcal{U}}^{\emptyset} * \vec{\mathcal{U}}^{\Delta_0^{\rightarrow}} * \cdots * \vec{\mathcal{U}}^{\Delta_n^{\rightarrow}}} \phi$. Suppose $\theta \mid\!\sim^{\Delta}_{lex} \phi$. If $\theta \models \bot$ then we are done, so suppose $\theta \not\models \bot$ and let $\langle i_0, \ldots, i_n \rangle$ be minimal under $\lhd_{lex}$ such that $\mathcal{U}_{i_0}^{\Delta_0^{\rightarrow}} \cap \ldots \cap \mathcal{U}_{i_n}^{\Delta_n^{\rightarrow}} \cap S_\theta \neq \emptyset$. Let $w_0 \in \mathcal{U}_{i_0}^{\Delta_0^{\rightarrow}} \cap \ldots \cap \mathcal{U}_{i_n}^{\Delta_n^{\rightarrow}} \cap S_\theta$. We must show $w_0 \in S_\phi$. Put $\Gamma = sent_{\Delta^{\rightarrow}}(w_0)$. Then $\Gamma \cup \{\theta\}$ is consistent. We claim that Γ is maximal under $\ll_{lex}$ with this property. To see this, note that, for each $r = 0, 1, \ldots, n$, $\Gamma_r = sent_{\Delta_r^{\rightarrow}}(w_0)$, and so $|\Gamma_r| = |sent_{\Delta_r^{\rightarrow}}(w_0)| = |\Delta_r| - i_r$, since $w_0 \in \mathcal{U}_{i_r}^{\Delta_r^{\rightarrow}}$. Now suppose there existed Γ' such that $\Gamma' \cup \{\theta\}$ was consistent and $\Gamma \ll_{lex} \Gamma'$, i.e., $\langle |\Delta_0| - i_0, \ldots, |\Delta_n| - i_n \rangle \lhd_{lex} \langle |\Gamma'_0|, \ldots, |\Gamma'_n| \rangle$. Then, by Lemma 3, this would give us $\langle |\Delta_0| - |\Gamma'_0|, \ldots, |\Delta_n| - |\Gamma'_n| \rangle \lhd_{lex} \langle i_0, \ldots, i_n \rangle$. Since $\Gamma' \cup \{\theta\}$ is consistent we know there exists $w'_0 \in W$ such that $w'_0 \models \bigwedge \Gamma' \wedge \theta$. Hence, for each $r = 0, 1, \ldots, n$, we have $|sent_{\Delta_r^{\rightarrow}}(w'_0)| \geq |\Gamma'_r|$ and so there exists $\langle j_0, \ldots, j_n \rangle \lhd_{lex} \langle i_0, \ldots, i_n \rangle$ such that $\mathcal{U}_{j_0}^{\Delta_0^{\rightarrow}} \cap \ldots \cap \mathcal{U}_{j_n}^{\Delta_n^{\rightarrow}} \cap S_\theta \neq \emptyset$ (take $j_r = |\Delta_r| - |sent_{\Delta_r^{\rightarrow}}(w'_0)|$ for each r). This contradicts the minimality of $\langle i_0, \ldots, i_n \rangle$ and so our claim is proved, i.e., Γ is indeed maximal. Hence, since $\theta \mid\!\sim^{\Delta}_{lex} \phi$ we have that $\Gamma \cup \{\theta\} \models \phi$ and so, since $w_0 \models \bigwedge \Gamma \wedge \theta$, we have $w_0 \models \phi$, i.e., $w_0 \in S_\phi$ as required.

For the converse direction of the theorem, suppose $\theta \mid\!\sim_{\vec{\mathcal{U}}^{\emptyset} * \vec{\mathcal{U}}^{\Delta_0^{\rightarrow}} * \cdots * \vec{\mathcal{U}}^{\Delta_n^{\rightarrow}}} \phi$. If $\theta \models \bot$ then for no $\Gamma \subseteq \Delta^{\rightarrow}$ do we have $\Gamma \cup \{\theta\}$ is consistent and so $\theta \mid\!\sim^{\Delta}_{lex} \phi$ holds vacuously. So suppose $\theta \not\models \bot$. Let Γ be a $\ll_{lex}$-maximal subset of $\Delta^{\rightarrow}$ such that $\Gamma \cup \{\theta\}$ is consistent and let $w_0 \in W$ be such that $w_0 \models \bigwedge \Gamma \wedge \theta$. We must show $w_0 \models \phi$. We have that $sent_{\Delta^{\rightarrow}}(w_0) \cup \{\theta\}$ is consistent and, for each r, $|sent_{\Delta_r^{\rightarrow}}(w_0)| \geq |\Gamma_r|$. But if, for some r, we had $|sent_{\Delta_r^{\rightarrow}}(w_0)| > |\Gamma_r|$, then we would have $\langle |\Gamma_0|, \ldots, |\Gamma_n| \rangle \lhd_{lex} \langle |sent_{\Delta_0^{\rightarrow}}(w_0)|, \ldots, |sent_{\Delta_n^{\rightarrow}}(w_0)| \rangle$, thus contradicting the $\ll_{lex}$-maximality of Γ. Hence we have that, for each r, $|sent_{\Delta_r^{\rightarrow}}(w_0)| = |\Gamma_r|$ and so $w_0 \in \mathcal{U}^{\Delta_r^{\rightarrow}}_{|\Delta_r| - |\Gamma_r|}$. Hence $\mathcal{U}^{\Delta_0^{\rightarrow}}_{|\Delta_0| - |\Gamma_0|} \cap \ldots \cap \mathcal{U}^{\Delta_n^{\rightarrow}}_{|\Delta_n| - |\Gamma_n|} \cap S_\theta \neq \emptyset$. If we can show that $\langle |\Delta_0| - |\Gamma_0|, \ldots, |\Delta_n| - |\Gamma_n| \rangle$ is minimal under $\lhd_{lex}$ such that this holds then we can use $\mathcal{U}^{\Delta_0^{\rightarrow}}_{|\Delta_0| - |\Gamma_0|} \cap \ldots \cap \mathcal{U}^{\Delta_0^{\rightarrow}}_{|\Delta_0| - |\Gamma_0|} \cap S_\theta \subseteq S_\phi$ (from $\theta \mid\!\sim_{\vec{\mathcal{U}}^{\emptyset} * \vec{\mathcal{U}}^{\Delta_0^{\rightarrow}} * \cdots * \vec{\mathcal{U}}^{\Delta_n^{\rightarrow}}} \phi$) to conclude that $w_0 \in S_\phi$ as required. But suppose

$\mathcal{U}^{\Delta_0^{\rightarrow}}_{j_0} \cap \ldots \cap \mathcal{U}^{\Delta_n^{\rightarrow}}_{j_n} \cap S_\theta \neq \emptyset$ for some $\langle j_0, \ldots, j_n \rangle \lhd_{lex} \langle |\Delta_0| - |\Gamma_0|, \ldots, |\Delta_n| - |\Gamma_n| \rangle$ and let $w'_0 \in \mathcal{U}^{\Delta_0^{\rightarrow}}_{j_0} \cap \ldots \cap \mathcal{U}^{\Delta_n^{\rightarrow}}_{j_n} \cap S_\theta$. Then $sent_{\Delta^{\rightarrow}}(w'_0) \cup \{\theta\}$ is consistent and $|sent_{\Delta_r^{\rightarrow}}(w'_0)| = |\Delta_r| - j_r$ for each $r = 0, 1, \ldots, n$. Thus, by Lemma 3, we have $\langle |\Gamma_0|, \ldots, |\Gamma_n| \rangle \lhd_{lex} \langle |sent_{\Delta_0^{\rightarrow}}(w'_0)|, \ldots, |sent_{\Delta_n^{\rightarrow}}(w'_0)| \rangle$ and so $\Gamma \ll_{lex} sent_{\Delta^{\rightarrow}}(w'_0)$. This contradicts the $\ll_{lex}$-maximality of Γ and so we must indeed have $\langle |\Delta_0| - |\Gamma_0|, \ldots, |\Delta_n| - |\Gamma_n| \rangle$ is minimal under $\lhd_{lex}$ such that $\mathcal{U}^{\Delta_0^{\rightarrow}}_{|\Delta_0|-|\Gamma_0|} \cap \ldots \cap \mathcal{U}^{\Delta_0^{\rightarrow}}_{|\Delta_0|-|\Gamma_0|} \cap S_\theta \neq \emptyset$ as required. ■

Using Propositions 8 and 9 we may re-express Theorem 2 as:

Corollary 1 *Let Δ be a set of defaults with associated Z-partition $(\Delta_0, \ldots, \Delta_n)$. Then, for all $\theta, \phi \in L$, we have $\theta \mid\!\sim^{\Delta}_{lex} \phi$ iff $\phi \in Bel(\preceq_\emptyset * \preceq_{\Delta_0^{\rightarrow}} * \cdots * \preceq_{\Delta_n^{\rightarrow}} * \preceq_\theta)$.*

If we go further and actually identify a revision of the form $\preceq * \preceq_E$ with $\preceq *E$ then we have the following characterisation of the lexicographic closure.

Corollary 2 *Let Δ be a set of defaults with associated Z-partition $(\Delta_0, \ldots, \Delta_n)$. Then, for all $\theta, \phi \in L$, we have $\theta \mid\!\sim^{\Delta}_{lex} \phi$ iff $\phi \in Bel(\preceq_\emptyset *\Delta_0^{\rightarrow} * \cdots * \Delta_n^{\rightarrow} * \theta)$.*

Hence, using *this* particular method of revision and *this* particular way of interpreting revision by a set of sentences, we have shown that $\theta \mid\!\sim^{\Delta}_{lex} \phi$ iff ϕ is believed after first successively revising the maximally ignorant epistemic state by the set of sentences $\Delta_i^{\rightarrow}$ for $i = 0, 1, \ldots, n$, and then revising by θ.[10]

5 The inclusion-based approach to default entailment

As we said in Section 1 the lexicographic closure of Δ is but one member of a family of consequence relations $\mid\!\sim^{\Delta}_{\ll}$, where $\ll$ is an ordering on $2^{\Delta^{\rightarrow}}$, and, for all $\theta, \phi \in L$, we have

$$\theta \mid\!\sim^{\Delta}_{\ll} \phi \quad \text{iff} \quad \begin{array}{l}\text{for all } \Gamma \subseteq \Delta^{\rightarrow} \text{ such that } \Gamma \cup \{\theta\} \text{ is consistent and } \Gamma \text{ is} \\ \ll\text{-maximal amongst such subsets, we have } \Gamma \cup \{\theta\} \models \phi,\end{array}$$

In this section we briefly describe another member of this family which was given by Brewka in [Bre 89]. Once again we base the ordering on the Z-partition $(\Delta_0, \ldots, \Delta_n)$[11]. The idea is to use the following *inclusion-based* ordering on

[10] We remark that this formulation of a sequence of revision steps in terms of a consequence relation is reminiscent of the *vertical* perspective of belief revision [Rot 96], according to which the operation of revision is reduced to simple addition of new information, without checking whether it is consistent with the prior beliefs. The actual beliefs of the agent at any one time are then retrieved from the hitherto collected information using some non-classical inference operator (in this case $\mid\!\sim^{\Delta}_{lex}$.)

[11] It should be pointed out that Brewka's approach, along with that of [BCDLP 93], is more general than the one in this paper, in that he considers the problem of generating a consequence relation from an arbitrary set of sentences partitioned according to some general notion of priority, rather than just concentrating on the Z-partition of a set of defaults

$2^{\Delta^{\rightarrow}}$, where once again, for $A, B \subseteq \Delta^{\rightarrow}$ we let $A_i = A \cap \Delta_i^{\rightarrow}$ for $i = 0, \ldots, n$ etc.,

$$\begin{array}{lll} A \ll_{ib} B & \text{iff} & \text{there exists } i \text{ such that } A_i \subset B_i \text{ and,} \\ & & \text{for all } j > i,\ A_j = B_j. \end{array}$$

Denoting the resulting consequence relation by $\mid\!\sim^{\Delta}_{ib}$, it is easy to see that $A \ll_{ib} B$ implies $A \ll_{lex} B$ and so, for all $\theta, \phi \in L$, we have $\theta \mid\!\sim^{\Delta}_{ib} \phi$ implies $\theta \mid\!\sim^{\Delta}_{lex} \phi$. Hence $\mid\!\sim^{\Delta}_{ib}$ represents a more *cautious* form of default entailment than $\mid\!\sim^{\Delta}_{lex}$. However, unlike $\mid\!\sim^{\Delta}_{lex}$, $\mid\!\sim^{\Delta}_{ib}$ will not, in general, be a rational consequence relation – it satisfies all the properties from Proposition 1 *except*, possibly, Rational Monotonicity (which makes it a (consistency-preserving) *preferential* consequence relation [KLM 90]). In [Neb 92] Nebel ties in $\mid\!\sim^{\Delta}_{ib}$ with his operator of *prioritized base revision*. A prioritized base is an arbitrary (not necessarily deductively closed) set of sentences P equipped with a total pre-ordering which reflects the relative *epistemic relevance* of its elements. When P is finite this pre-ordering may be represented as a partition $(P_0, \ldots, P_m)$ of P, where P_m is the set of the most epistemically relevant elements of P and P_0 is the set of the least epistemically relevant elements. The prioritized base revision operator $\hat{\oplus}$ takes as arguments a prioritized base and a sentence representing the epistemic input and returns a belief set. According to the correspondence shown by Nebel, we have

$$\theta \mid\!\sim^{\Delta}_{ib} \phi \text{ iff } \phi \in (\Delta_0^{\rightarrow}, \ldots, \Delta_n^{\rightarrow}) \,\hat{\oplus}\, \theta$$

This formula bears comparison to the one we obtained in Corollary 2, suggesting a close connection between the revision method studied in this paper and prioritized base revision. We leave working out the precise details of this connection to another occasion. However we remark here that $\hat{\oplus}$ and $*$ do still differ in two important regards: (1) $\hat{\oplus}$ takes only single sentences as epistemic inputs whereas $*$ accepts more general inputs, in particular sets of sentences, and (2) $\hat{\oplus}$ returns as output only a new belief set whereas $*$ returns not just a new belief set but also the structure (in the form of a new E-relation) required to carry out further revisions. Thus, as it stands, $\hat{\oplus}$ is, in contrast to $*$, incapable of iterated revision.

6 Further work

The developments in the previous sections have raised a couple of questions regarding both belief revision *and* default entailment. Firstly, while there have been several papers published concerned with iterated revision by single sentences, and also some concerned with revision by sets of sentences,[12] there seems to be little in the way of any systematic study of iterated revision *by* sets

[12] Either directly (e.g. [Zha 96]) or indirectly, via the study of *contraction* by a set of sentences (e.g. [FH 94]). See [Gär 88] for a description of contractions and their close relationship with revision.

of sentences.[13] Darwiche and Pearl [DP 97] provide a postulational approach to the question of iterated revision of epistemic states by single sentences. In this approach they take the concept of epistemic state to be primitive, assuming only that from each such state Ψ we may extract a belief set (in the usual AGM sense of the term) $B(\Psi)$ representing the set of sentences accepted in that state. For example Darwiche and Pearl's second postulate may be stated as

$$\text{If } \phi \models \neg\theta \text{ then } B((\Psi * \theta) * \phi) = B(\Psi * \phi).$$

(For the other postulates and their justifications see [DP 97].) It is not difficult to see that, if we identify epistemic states with E-relations and take $B(\preceq) = Bel(\preceq)$, then the method proposed by Nayak, on its restriction to single sentences[14] satisfies all of Darwiche and Pearl's postulates. However, it also satisfies some interesting properties in the general case. For example, given an E-relation $\preceq$ and $E_1 \subseteq E_2 \subseteq L$ such that E_2 is consistent, it can be shown that $(\preceq *E_2) * E_1 = (\preceq *E_2 - E_1) * E_1$. In particular, if $\{\theta, \phi\}$ is consistent, we have $(\preceq *\{\theta, \phi\}) * \phi = (\preceq *\theta) * \phi$. (Note this is a stronger statement than just $Bel((\preceq *\{\theta, \phi\}) * \phi) = Bel((\preceq *\theta) * \phi)$.) The question of whether this, or any other, property of iterated revision by sets is desirable seems to be a question worth investigating. Another question is: Can we, by modifying the various parameters involved in this revision process, model any of the other existing methods of default entailment, apart from the lexicographic closure, or even construct new ones? For example, given our set of defaults Δ and its Z-partition $(\Delta_0, \ldots, \Delta_n)$, let $\Theta_i = \bigcup_{i \leq j} \Delta_j$ for each $i = 1, \ldots, n$. Then, by the above comments, we may rewrite Corollary 2 as

$$\theta \mid\!\sim^{\Delta}_{lex} \phi \text{ iff } \phi \in Bel(\preceq_{\emptyset} *\Theta_0^{\rightarrow} * \cdots * \Theta_n^{\rightarrow} * \theta).$$

We conjecture that if we now replace each $\Theta_i^{\rightarrow}$ in the above by $\bigwedge \Theta_i^{\rightarrow}$, then we obtain the *rational closure* [LM 92] (which is semantically equivalent to System Z [Pea 90]) of Δ, instead of the lexicographic closure. This and other variations are the subject of ongoing study. Finally, note that, since we assumed at the outset that our language L is based on only finitely many propositional variables, and also that Δ is a finite set of defaults, we have not needed in this paper to confront the question of revision by *infinite* sets of sentences. It remains to be seen to what extent the ideas in this paper can be extended to cover this more general situation.[15]

Conclusion

In this paper we have taken a particular model of default reasoning – the lexicographic closure – and re-cast it in terms of iterated belief revision by sets

[13] An exception, in a slightly more complex framework, is [Wey 99].

[14] We obviously interpret single sentences here as singleton sets.

[15] For one treatment of this topic, and its relation with nonmonotonic inference from infinite sets of premises, see [ZCZL 97].

of sentences, using the particular, independently motivated, revision model of Nayak. In the process of doing this, a couple of interesting avenues for further exploration have suggested themselves. In particular, the questions of which properties of iterated multiple revision should be deemed desirable, and of how we may apply the principles underlying the AGM theory of belief revision in the context of default reasoning.

Acknowledgements

Much of this paper was written while the author was a researcher at the Max-Planck-Institute for Computer Science in Saarbrücken, Germany. The author would like to thank Emil Weydert, Abhaya Nayak, Michael Freund, Hans Rott, Leon van der Torre and the anonymous referees of both the NMR2000 workshop and this journal for helpful comments and suggestions.

References

[AGM 85] C. E. Alchourŕon, P. Gärdenfors and D. Makinson, On the logic of theory change: partial meet functions for contraction and revision, *Journal of Symbolic Logic* **50** (1985) 510–530.

[BCDLP 93] S. Benferhat, C. Cayrol, D. Dubois, J. Lang and H. Prade, Inconsistency management and prioritized syntax-based entailment, in: R. Bajcsy, ed, *Proceedings International Joint Conference on Artificial Intelligence*, (Morgan Kaufmann, 1993) 640–645.

[BSS 95] S. Benferhat, A. Saffiotti and P. Smets, Belief functions and default reasoning, in: *Proceedings of the 11th Uncertainty in A.I. Conference*, (Morgan Kaufmann, 1995) 19–26.

[Bou 96] C. Boutilier, Iterated revision and minimal change of conditional beliefs, *Journal of Philosophical Logic* **25** (1996) 262–305.

[Bre 89] G. Brewka, Preferred subtheories: An extended logical framework for default reasoning, in: N. S. Sridharan (ed.), *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, (Morgan Kaufmann, 1989) 1043–1048.

[DP 97] A. Darwiche and J. Pearl, On the logic of iterated belief revision, *Artificial Intelligence* **89** (1997) 1–29.

[Fre 93] M. Freund, Injective models and disjunctive relations, *Journal of Logic and Computation* **3** (3) (1993) 231–247.

[Fre 99] M. Freund, Statics and dynamics of induced systems, *Artificial Intelligence* **110** (1) (1999) 103–134.

[FH 99] N. Friedman and J. Y. Halpern, Belief revision: A critique, *Journal of Logic, Language and Information* **8** (4) (1999) 401–420.

[FH 94] A. Fuhrmann and S. O. Hansson, A survey of multiple contractions, *Journal of Logic, Language and Information* **3** (1994) 39–76.

[Gär 88] P. Gärdenfors, *Knowledge in Flux*, MIT Press, (1988).

[GM 94] P. Gärdenfors and D. Makinson, Nonmonotonic inferences based on expectations, *Artificial Intelligence* **65** (2) (1994) 197–245.

[GR 95] P. Gärdenfors and H. Rott, Belief revision, in: D. M. Gabbay, C. J. Hogger and J. A. Robinson (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 4: Epistemic and Temporal Reasoning*, Oxford University Press, (1995) 35–132.

[GMP 93] M. Goldszmidt, P. Morris and J. Pearl, A maximum entropy approach to nonmonotonic reasoning, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **15** (1993) 220–232.

[KLM 90] S. Kraus, D. Lehmann and M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, *Artificial Intelligence* **44** (1990) 167–207.

[Leh 95] D. Lehmann, Another perspective on default reasoning, *Annals of Mathematics and Artificial Intelligence* **15** (1995) 61–82.

[LM 92] D. Lehmann and M. Magidor, What does a conditional knowledge base entail?, *Artificial Intelligence* **55** (1992) 1–60.

[Mak 94] D. Makinson, General patterns in nonmonotonic reasoning, in: D. M. Gabbay, C. J. Hogger and J. A. Robinson (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*, Oxford University Press, (1994) 35–110.

[MG 90] D. Makinson and P. Gärdenfors, Relations between the logic of theory change and nonmonotonic logic, in: A. Fuhrmann, M. Morreau (eds.), *The Logic of Theory Change*, LNAI 465, (Springer Verlag, 1990) 185–205.

[NNP 96] A. C. Nayak, P. Nelson and H. Polansky, Belief change as change in epistemic entrenchment, *Synthese* **109**(2) (1996) 143–174.

[Nay 94] A. C. Nayak, Iterated belief change based on epistemic entrenchment, *Erkenntnis* **41** (1994) 353–390.

[Neb 92] B. Nebel, Syntax-based approaches to belief revision, in: P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, (1992) 52–88.

[Pea 90] J. Pearl, System Z: a natural ordering of defaults with tractable applications to default reasoning, in: *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, (Morgan Kaufmann, San Mateo, CA, 1990) 121–135.

[Rot 92a] H. Rott, Modellings for belief change: prioritization and entrenchment, *Theoria* **58** (1992) 21–57.

[Rot 92b] H. Rott, Preferential belief change using generalized epistemic entrenchment, *Journal of Logic, Language and Information* **1** (1992) 45–78.

[Rot 96] H. Rott, *Making Up One's Mind: Foundations, Coherence, Nonmonotonicity*, Habilitationsschrift, University of Konstanz (1996) (revised version to appear with Oxford University Press, new title *Change, Choice and Inference*).

[Spo 88] W. Spohn, Ordinal conditional functions: a dynamic theory of epistemic states, in: W. Harper and B. Skyrms (eds.), *Causation in Decision, Belief Change and Statistics* (Dordrecht, Netherlands: Reidel 1988) Volume 2, 105–134.

[Wey 96] E. Weydert, System J - revision entailment, in *Proceedings of the International Conference on Formal and Applied Practical Reasoning (FAPR'96)* (1996)

[Wey 99] E. Weydert, JZBR - Iterated belief change for conditional ranking constraints, in: *Proceedings of the 4th Dutch-German Workshop on Nonmonotonic Reasoning Techniques and Their Applications (DGNMR'99)*, (ILLC Scientific Publications, Amsterdam, 1999).

[Wil 94] M. A. Williams, Transmutations of Knowledge Systems, in: J. Doyle and E. Sandewall (eds.), *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR'94)*, (Morgan Kaufmann, 1994) 619–629.

[Zha 96] D. Zhang, Belief revision by sets of sentences, *Journal of Computer Science and Technology* **11**(2) (1996) 1–19.

[ZCZL 97] D. Zhang, S. Chen, W. Zhu and H. Li, Nonmonotonic reasoning and multiple belief revision, in: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'97)*, (Morgan Kaufmann, 1997) 95–100.