

Belief Revision, Conditional Logic and Nonmonotonic Reasoning

WAYNE WOBCKE

Abstract We consider the connections between belief revision, conditional logic and nonmonotonic reasoning, using as a foundation the approach to theory change developed by Alchourrón, Gärdenfors and Makinson (the AGM approach). This is first generalized to allow the iteration of theory change operations to capture the dynamics of epistemic states according to a principle of minimal change of entrenchment. The iterative operations of expansion, contraction and revision are characterized both by a set of postulates and by Grove's construction based on total pre-orders on the set of complete theories of the belief logic. We present a sound and complete conditional logic whose semantics is based on our iterative revision operation, but which avoids Gärdenfors's triviality result because of a severely restricted language of beliefs and hence the weakened scope of our extended postulates. In the second part of the paper, we develop a computational approach to theory dynamics using Rott's E-bases as a representation for epistemic states. Under this approach, a ranked E-base is interpreted as standing for the most conservative entrenchment compatible with the base, reflecting a kind of foundationalism in the acceptance of evidence for a belief. Algorithms for the computation of our iterative versions of expansion, contraction and revision are presented. Finally, we consider the relationship between nonmonotonic reasoning and both conditional logic and belief revision. Adapting the approach of Delgrande, we show that the unique extension of a default theory expressed in our conditional logic of belief revision corresponds to the most conservative belief state which respects the theory: however, this correspondence is limited to propositional default theories. Considering first order default theories, we present a belief revision algorithm which incorporates the assumption of independence of default instances and propose the use of a base logic for default reasoning which incorporates uniqueness of names. We conclude with an examination of the behavior of an implemented system on some of Lifschitz's benchmark problems in nonmonotonic reasoning.

1 Introduction Much recent work in belief revision uses mathematical methods, building on the work of Alchourrón, Gärdenfors and Makinson [1]. In this approach, known as the AGM approach, a belief state is modeled as a logical theory over a base logic constrained only to be consistent and compact, contain the Propositional Cal-

culus and *modus ponens* and satisfy the deduction theorem. Three basic operations on belief states are defined: *expansion* is the addition to a belief set of a belief and its consequences; *contraction* is the removal of a belief from a belief set; and *revision* is the incorporation of a belief into a belief set with which it may be inconsistent. But not any function from belief sets to belief sets counts as a *rational* belief change operation. The idea is to capture some aspects of rationality by specifying postulates that such belief change functions should satisfy. One supposedly rational property of belief change operations is minimal change, i.e., when a belief set has to be modified, only those changes are made which are necessary to ensure the operation's success. A central concern is to make this notion of minimal change mathematically precise.

The notion of minimal change has a long history in belief revision, dating at least from Ramsey [55], who introduced, but did not formalize, the Ramsey test for accepting a conditional statement. The Ramsey test was adapted by Stalnaker [59] and further by Lewis [40] in providing possible worlds semantics for counterfactual conditionals. This has resulted in mathematical models of minimality in change. In Artificial Intelligence, the intuition of minimal change figures in the modeling of actions in that when computing the effect of an action, it is assumed that as little as necessary changes in the agent's model of the world, McCarthy and Hayes [45], Ginsberg and Smith [28]. Thus minimal change, in either its semantic or epistemic forms, is a widely-held intuition underlying the dynamics of a rational agent's mental states.

Any logical formulation of belief revision, the semantics of conditionals and reasoning about action must fail to satisfy monotonicity: if $\Gamma \vdash A$ then $\Gamma' \vdash A$ whenever $\Gamma \subseteq \Gamma'$. Nonmonotonic reasoning has been studied extensively in Artificial Intelligence and various formalisms exist for representing and reasoning with nonmonotonicity. This work grew, in part, out of attempts to formalize reasoning with inheritance networks—network structures with a procedural inference system that can be used to represent defaults such as that birds can typically fly and their exceptions such as that penguins cannot typically fly (even though all penguins are birds). Exceptions can also have exceptions, etc. Nonmonotonicity arises when considering the new information that Tweety is a bird, from which it follows (intuitively) that Tweety flies, although this conclusion will have to be retracted on acquiring the additional information that Tweety is a penguin. An important point is that although inference by default requires nonmonotonicity, it is not only nonmonotonicity that has to be handled: solving problems in nonmonotonic reasoning also requires methods for dealing with partial, or incomplete, initial information, as emphasized by Etherington [15].

There is an obvious intuitive connection between conditional logic and nonmonotonic reasoning. The default inference of Tweety's flying from Tweety's being a bird corresponds closely to the indicative conditional "if Tweety is a bird then Tweety flies." Indeed this correspondence has been investigated formally by Delgrande [9]. Due to the already identified connection between conditional logic and belief revision underpinned by the principle of minimal change, there is strong evidence for close links between all three areas of inquiry: belief revision, conditional logic and nonmonotonic reasoning. It is the purpose of this paper to consider these connections more formally.

Our work starts with a generalization of the AGM approach to belief revision. It is a widely-accepted weakness of the AGM approach that belief change operations

are defined with respect to belief sets—the *contents* of belief states—rather than with respect to belief states themselves. Thus AGM minimal change is a very weak notion. The problem of representing both the contents and dynamics of belief states has been addressed by Gärdenfors and Makinson [25] who proposed epistemic entrenchments as a means of representing belief states so that particular belief change operations could also be represented (epistemic entrenchment is the epistemic analogue of what Lewis [40] calls comparative possibility). We first propose a set of extended AGM postulates for epistemic state change, based on interpreting the notion of minimal change as applying to belief states represented as entrenchments over belief sets. The iterative theory change operations of expansion, contraction and revision are characterized using sets of postulates and using Grove’s [32] construction based on total pre-orders on the set of complete theories of the base logic.

Then by adapting methods familiar from conditional logic, we define a conditional logic of belief revision that captures our version of minimal change. However, the problem of defining revision by conditionals is not considered in this work, i.e., under our approach to the revision of epistemic states, the new information to be accepted must be nonconditional information. Moreover, the only conditional information believed by the agent must be derived from entrenchments over the base language and this base language cannot contain conditional operators. Thus we avoid Gärdenfors’s [19] triviality result by restricting an agent’s “freedom” to believe in conditionals. Essentially, the triviality result states that with a base logic containing conditional operators, any belief system in which revisions are in accord with the AGM postulates and which satisfy the Ramsey test (i.e., $A \Rightarrow B$ is accepted in a state iff B is accepted in the state resulting by a revision to accept A) must be in some sense trivial. The problem stems, in part, from the application of the AGM postulates to conditional beliefs, and we simply rule out this possibility.

Another motivation for our restrictive approach to dealing with the Ramsey test is our goal of developing a computational version of belief revision able to be used for solving problems in nonmonotonic reasoning. The AGM approach is not computational because epistemic entrenchments (the representations of belief states) are total pre-orders on the set of formulas over the base logic of beliefs. Some of these orderings are not finitely representable: there may be infinitely many different levels of entrenchment. Following Rott [56], we use an E-base to represent an epistemic state: a ranked E-base is understood to stand for a uniquely determined entrenchment which is the most “conservative” entrenchment which extends it—rather as a partially specified theory is understood as standing for its logical closure, the smallest closed theory containing it. Our approach to partial information is indeed conservative. Using the work of Williams [61], we consider the special case of a ranked entrenchment as providing for each belief a natural number known as its rank, which reflects the degree of evidence for the corresponding belief. Given an incomplete entrenchment and a formula whose rank is not explicitly known, the available evidence for that formula consists of all the proofs of the formula from the belief set: the unknown rank must be at least that of each minimally ranked formula used in such a proof. We are conservative in that we take the unknown rank to be exactly the maximum of the ranks of those formulas which are minimally ranked in a proof of the formula. To substantiate the claim that this approach is computational, we present algorithms for the computation

of expansion, contraction, and revision according to our theory of minimal change of entrenchment.

Incompletely determined entrenchments are not merely of computational concern: incomplete information is a key aspect of problems in nonmonotonic reasoning, so methods for dealing with partial information about entrenchments form a central part of this paper. Our idea of conservatism embodies a strong element of foundationalism, in contrast to the purely coherence approach to epistemics advocated by Gärdenfors [21]. However, we claim that at least this much foundationalism is essential to handling problems in default reasoning. In the final part of this paper, we consider the connections between nonmonotonic reasoning and both conditional logic and belief revision. In the case of propositional theories, adapting Delgrande's [10] approach, we show that the unique extension of a default theory expressed in conditional logic is exactly the most conservative belief state that respects the theory. In the case of first order logic with an equality predicate, we define a special revision operation appropriate for nonmonotonic reasoning, which in addition to conservatism in accepting evidence, incorporates the assumption of independence of default instances. Our belief revision approach to nonmonotonic reasoning has been implemented in a computer system, the details of which are described in Dixon and Wobcke [12]: the heart of the system is a standard resolution theorem prover incorporating uniqueness of names, an assumption which is supported by intuitions in default reasoning. We conclude with an examination of the behavior of our system on some benchmark problems in nonmonotonic reasoning collected by Lifschitz [41].

2 The AGM approach to theory change

2.1 Expansions, contractions, revisions In the AGM approach to formalizing belief change, belief states are modeled as logically closed sets of formulas over a base logic constrained only to be consistent and compact, contain the Propositional Calculus and *modus ponens* and satisfy the deduction theorem. Three operations are defined: the *expansion* of the belief set K by A , denoted K_A^+ , represents the addition of A to K ; the *contraction* of K by A , denoted K_A^- , represents the removal of A from K ; the *revision* of K by A , denoted K_A^* , represents the addition of A to K so as to preserve consistency. In a revision, if the new belief A is inconsistent with K , some of the original beliefs must be given up in order to accept A , whereas in an expansion, A is added to K regardless of whether it is inconsistent with K . Alchourrón, Gärdenfors and Makinson [1] formulated the following "rationality postulates" which state the desired properties of the belief sets resulting from each of the theory change operations. In the following definitions, Cn is the logical consequence operation for the base logic, and K_\perp is the inconsistent set of all beliefs.

First, six postulates for expansion are presented.

- (**K⁺1**) K_A^+ is a belief set.
- (**K⁺2**) $A \in K_A^+$.
- (**K⁺3**) $K \subseteq K_A^+$.
- (**K⁺4**) If $A \in K$ then $K_A^+ \subseteq K$.
- (**K⁺5**) If $K \subseteq H$ then $K_A^+ \subseteq H_A^+$.
- (**K⁺6**) K_A^+ is the smallest belief set satisfying (**K⁺1**)–(**K⁺5**).

These postulates have the consequence that $K_A^+ = Cn(K \cup \{A\})$.

The following postulates for contraction are given.

- (**K⁻1**) K_A^- is a belief set.
- (**K⁻2**) $K_A^- \subseteq K$.
- (**K⁻3**) If $A \notin K$ then $K_A^- = K$.
- (**K⁻4**) If $\not\vdash A$ then $A \notin K_A^-$.
- (**K⁻5**) If $A \in K$ then $K \subseteq (K_A^-)^+$.
- (**K⁻6**) If $\vdash A \leftrightarrow B$ then $K_A^- = K_B^-$.
- (**K⁻7**) $K_A^- \cap K_B^- \subseteq K_{A \wedge B}^-$.
- (**K⁻8**) If $A \notin K_{A \wedge B}^-$ then $K_{A \wedge B}^- \subseteq K_A^-$.

Finally, the following postulates for revision are defined.

- (**K^{*}1**) K_A^* is a belief set.
- (**K^{*}2**) $A \in K_A^*$.
- (**K^{*}3**) $K_A^* \subseteq K_A^+$.
- (**K^{*}4**) If $\neg A \notin K$ then $K_A^+ \subseteq K_A^*$.
- (**K^{*}5**) $K_A^* = K_\perp$ only if $\vdash \neg A$.
- (**K^{*}6**) If $\vdash A \leftrightarrow B$ then $K_A^* = K_B^*$.
- (**K^{*}7**) $K_{A \wedge B}^* \subseteq (K_A^*)_B^+$.
- (**K^{*}8**) If $\neg B \notin K_A^*$ then $(K_A^*)_B^+ \subseteq K_{A \wedge B}^*$.

These postulates are motivated on the grounds that accepting a new belief A should result in a “minimal” disturbance to the set K of existing beliefs. For example, (**K^{*}4**) says that if A is consistent with K then no beliefs from K need be removed in accepting A , and (**K^{*}3**) says that in any case, there cannot be extra beliefs in the revised set K_A^* that do not follow from K and A . (**K^{*}7**) and (**K^{*}8**) are the analogues of these for acceptance of conjunctive information. So (**K^{*}8**) says that if, after accepting A , B is consistent, then no beliefs from K_A^* need be removed in accepting $A \wedge B$, and (**K^{*}7**) says that no extra beliefs can be included in $K_{A \wedge B}^*$ that do not follow from K_A^* and B . An alternative interpretation of this notion of minimal change can be based on the idea of partial meet contraction (see [1]), in which the set K_A^- is defined as the intersection of a favored subset of the maximal subsets of K that do not contain A .

The contraction and revision postulates are not constructive in the sense that there are many functions that satisfy the postulates and so are “rational” belief change operations. However, it can be shown that particular contraction and revision operations can be related to each other by the Levi and Harper identities. That is, given any contraction operation satisfying (**K⁻1**)–(**K⁻8**), a revision operation satisfying (**K^{*}1**)–(**K^{*}8**) can be defined using the Levi identity: $K_A^* = (K_{\neg A}^-)^+$. Similarly, given any revision operation satisfying (**K^{*}1**)–(**K^{*}8**), a contraction operation satisfying (**K⁻1**)–(**K⁻8**) can be defined using the Harper identity: $K_A^- = K \cap K_{\neg A}^*$. Moreover, the Levi and Harper identities are duals to each other in the sense that the revision operation obtained using the Levi identity from the contraction operation defined using the Harper identity based on some revision operation is the same as the original revision operation, and similarly, starting with a contraction operation and using first the Levi, then the Harper, identity results in a contraction operation identical to the original operation.

It is important to note that in the AGM framework, belief states can be sets of formulas over any base logic satisfying the above conditions. Typically, the base logic is a classical logic, but the base logic could be a modal or conditional logic. When conditional logic is used as the base logic, there arises a special problem. The desire is to interpret the conditional operator \Rightarrow as signifying revision of beliefs, so that belief in conditionals is interpreted as conditional belief. The belief of a conditional $A \Rightarrow B$ (in an initial state) is to be interpreted as the belief of B conditional on A (in that state), which amounts to full belief in B in the state resulting from the initial state after a revision to accept A has occurred. Naturally this interpretation makes sense if and only if the beliefs in conditionals in fact reflect the actual state of affairs, i.e., a belief $A \Rightarrow B$ is held iff B is accepted after a revision to accept A . This is the Ramsey test for the truth of the conditional, from [55].

Ramsey, of course, did not have in mind Gärdenfors's particular postulates for belief revision when he formulated his test. Thus the question of whether the Ramsey test can be fulfilled under these conditions on revision functions is a real one. Perhaps surprisingly, given the strong intuitions behind both the Ramsey test and the AGM postulates, the answer is that they are incompatible. This is the Gärdenfors triviality result, reported in [19].

Definition 2.1 A *belief set* is a logically closed set of formulas over a base logic.

Definition 2.2 A *belief system* is a collection of belief sets together with a revision function $*$ from belief sets and formulas to belief sets that satisfies **(K*1)**–**(K*8)**.

Definition 2.3 A belief system satisfies the Ramsey test if for all belief sets K and all formulas A and B , K contains $A \Rightarrow B$ iff K_A^* contains B .

Definition 2.4 Two formulas A and B are *disjoint* if $\vdash \neg(A \wedge B)$.

Definition 2.5 A belief system is *nontrivial* if there are at least three pairwise disjoint sentences A , B and C and some belief set K in the system that is consistent with all three sentences, i.e., $\neg A \notin K$, $\neg B \notin K$ and $\neg C \notin K$.

To incorporate conditional belief and the Ramsey test into the theory, the base logic must be a conditional logic, but which one? To answer this question, Gärdenfors [19] defines the following notion of validity, and takes the base logic **C** to be the set of all such valid sentences.

Definition 2.6 A sentence A is *valid* if in any belief system, the only belief set containing $\neg A$ is the inconsistent belief set.

The logic **C** consists of the following axiom schemata and inference rules.

- (A1) All truth-functional tautologies.
- (A2) $(A \Rightarrow B) \wedge (A \Rightarrow C) \rightarrow (A \Rightarrow (B \wedge C))$.
- (A3) $A \Rightarrow \text{true}$.
- (A4) $A \Rightarrow A$.
- (A5) $(A \Rightarrow B) \rightarrow (A \rightarrow B)$.
- (A6) $(A \wedge B) \rightarrow (A \Rightarrow B)$.
- (A7) $(A \Rightarrow \neg A) \rightarrow (B \Rightarrow \neg A)$.
- (A8) $(A \Rightarrow B) \wedge (B \Rightarrow A) \rightarrow ((A \Rightarrow C) \rightarrow (B \Rightarrow C))$.
- (A9) $(A \Rightarrow C) \wedge (B \Rightarrow C) \rightarrow ((A \vee B) \Rightarrow C)$.

- (A10) $(A \Rightarrow B) \wedge \neg(A \Rightarrow \neg C) \rightarrow ((A \wedge C) \Rightarrow B)$.
 (MP) From A and $A \rightarrow B$ infer B .
 (RCM) If $\vdash B \rightarrow C$ infer $(A \Rightarrow B) \rightarrow (A \Rightarrow C)$.

There is a loose (but not exact) correspondence between these axiom schemata and the postulates. But importantly not all the postulates are needed to derive the triviality result, in which case not all the above axiom schemata are needed in the base logic. In fact, the triviality result can be proven without reference to the base logic, using the following property of belief sets (monotonicity) which is a consequence of the Ramsey test.

- (K*M) If $K \subseteq H$ then $K_A^* \subseteq H_A^*$.

Theorem 2.7 ([19]) *Any belief system over the base logic corresponding to (K*2), (K*4), (K*5) and satisfying (K*M) is trivial.*

This result motivates our alternative modeling of belief revision based on a theory of the dynamics of epistemic states. We use this theory as a semantics for a conditional logic: a conditional $A \Rightarrow B$ holds at a belief state K if B holds at the state resulting from the revision of K to accept A with some degree of strength, as determined by a selection function. The revision operation must satisfy a set of generalized AGM postulates designed to capture the principle of minimal change of entrenchment. We give a sound and complete axiomatization of our conditional logic. The triviality result is avoided because of the weakened language of beliefs and the restricted “freedom” of an agent to have beliefs in conditionals. Essentially, the only conditional beliefs an agent can have are those that could result from a sequence of revisions each obeying the principle of minimal change. The net effect of using only certain of the possible revision functions to represent the possible conditional beliefs is that the full force of the AGM postulates cannot be applied to conditional beliefs. Note also that the inconsistent belief state requires special treatment because the condition of logical closure implies that the corresponding set of beliefs contains all formulas of the language, in particular $A \Rightarrow B$ and $A \Rightarrow \neg B$ for any A and B . Hence if the Ramsey test is to be fulfilled when $K = K_\perp$, K_A^* would have to be K_\perp for any A . But this is in direct conflict with (K*5).

2.2 Epistemic entrenchment For a given collection of belief sets, there are typically many contraction and revision operations satisfying the AGM postulates. The reason for this is that, although the postulates are an attempt to capture a notion of minimal change, no purely logical definition of minimal change is possible: the notion of minimality used will be dependent on the domain or context of the belief system, just as Lewis’s [40] notion of comparative similarity in possible worlds semantics for conditionals is also context dependent. Thus a method for representing particular contraction and revision operations is necessary and is of interest both mathematically and computationally.

Gärdenfors and Makinson [25] introduce the use of epistemic entrenchments to represent particular contraction and revision operations. An *epistemic entrenchment* is a total pre-order \leq (with strict part $<$) on formulas of the base logic, intuitively representing the degree of strength with which a belief is held. An ordering is an epistemic entrenchment if it satisfies the following five conditions.

- (EE1) If $A \leq B$ and $B \leq C$ then $A \leq C$.
- (EE2) If $A \vdash B$ then $A \leq B$.
- (EE3) For any A and B , $A \leq A \wedge B$ or $B \leq A \wedge B$.
- (EE4) When $K \neq K_{\perp}$, $A \notin K$ iff $A \leq B$ for all B .
- (EE5) If $B \leq A$ for all B , then $\vdash A$.

Contractions and revisions are defined using the (C^-) and (C^*) conditions. It is shown that every contraction and revision operation defined as using epistemic entrenchment satisfies the AGM postulates, and conversely, every contraction and revision operation satisfying the AGM postulates can be represented by an epistemic entrenchment.

- (C^-) $B \in K_A^-$ iff $B \in K$ and either $\vdash A$ or $A < A \vee B$.
- (C^*) $B \in K_A^*$ iff either $\vdash \neg A$ or $\neg A < A \rightarrow B$.

Epistemic entrenchments are closely related to the comparative possibility orderings defined by Lewis [40] in his analysis of counterfactual conditionals. In both cases, the orderings are total pre-orders. The main difference is that for an epistemic entrenchment, the least entrenched formulas are the nonbeliefs, whereas for comparative possibility the formulas true in the actual world are minimal in the ordering. In fact, given a comparative possibility relation \leq_{cp} on a set of complete theories (the epistemic analogue of possible worlds), an epistemic entrenchment \leq_{ee} can be defined by $A \leq_{ee} B$ iff $\neg A \leq_{cp} \neg B$, and vice versa.

Given the close correspondence between epistemic entrenchment and comparative possibility, it is unsurprising that a model-theoretic characterization of belief revision operations can be based on systems of spheres of complete theories. Such a characterization was provided by Grove [32]. The main difference here is that whereas Lewis's spheres are centered on the actual world, Grove's spheres are centered on the set of complete theories containing the belief set. Grove's definitions are repeated below: they will be adapted in the next section dealing with the dynamics of epistemic states.

Let $[K]$ denote the set of consistent complete theories containing K .

Definition 2.8 A *system of spheres centered on* $[K]$ is a collection \mathbf{S} of sets of consistent complete theories that satisfies the following conditions.

- (S1) \mathbf{S} is totally ordered by \subseteq , i.e., if S and S' are in \mathbf{S} then $S \subseteq S'$ or $S' \subseteq S$.
- (S2) $[K]$ is the \subseteq -minimum of \mathbf{S} , i.e., $[K] \in \mathbf{S}$ and if $S \in \mathbf{S}$ then $[K] \subseteq S$.
- (S3) The set of all consistent theories is an element of \mathbf{S} .
- (S4) If A is a formula and some sphere in \mathbf{S} intersects $\{A\}$ then there is a smallest sphere in \mathbf{S} that intersects $\{A\}$.

A revision operation can be defined for systems of spheres just as Lewis defines the truth conditions for conditionals: given a consistent formula A , let S_A be the smallest sphere intersecting $\{A\}$. Then K_A^* is the intersection of the theories contained in $\min_{\mathbf{S}}(A) = S_A \cap \{A\}$ (if no theory contains A then K_A^* is K_{\perp}). It is shown that revision operations so defined satisfy all the AGM postulates $(\mathbf{K}^*1) - (\mathbf{K}^*8)$, and conversely, any revision operation can be modeled as a system of spheres so that for each consistent formula A , K_A^* is the intersection of theories contained in $\min_{\mathbf{S}}(A)$. Similar results hold for contractions when K_A^- is defined as the intersection of theories

contained in $[K] \cup \text{min}_S(\neg A)$. An illustration of a system of spheres is presented in the next section.

3 Dynamics of epistemic states It is a widely-recognized weakness of the AGM approach to theory change that iteration of belief change operations is impossible. If a belief state is a set of beliefs together with several possible belief change operations, then this is because each such operation is defined as mapping a belief set and a formula to a belief set, whereas for an operation to be iterative, the result should be another belief *state*. Equivalently, if epistemic entrenchments are used to specify theory change operations, there is no way to specify the entrenchment on the set of beliefs resulting from a belief change operation.

There are a number of ways to address to this problem, each of which has received considerable attention. Historically, the first was to extend the base logic to include conditional formulas as in Gärdenfors [19]: then using the (C^-) and (C^*) conditions, conditional beliefs could hold in the resulting states just as nonconditional beliefs. As described above, this leads immediately to the triviality result. Either the Ramsey test or the AGM postulates (or both) have to be weakened. If it is desired to retain the AGM postulates, one might first try to weaken the Ramsey test. However, various proposals, e.g., by Gärdenfors [20] and by Gärdenfors, Lindström, Morreau and Rabinowicz [24] have all failed to avoid the triviality result.

Another way to avoid the triviality result is to insist on an epistemological difference between nonconditional and conditional beliefs, requiring that the Ramsey test apply only to nonconditional beliefs. Levi [39] is a major proponent of this position, and in an AGM context, this line is taken by Rott [56] and Morreau [47]. In such an approach, belief states are logically closed nonconditional theories, with conditional beliefs derived from the action of a revision operation. The trouble with this kind of general framework is that with a natural way of defining the semantics of conditional beliefs such as that in Wobcke [63], none of the interesting AGM postulates (**K*3**), (**K*4**), (**K*7**) and (**K*8**) are valid in an unrestricted form. Thus if these postulates are to be retained in some form, restrictions have to be placed on the class of epistemic states and possibly also on the class of allowable revision operations.

In a non-AGM context, an approach to modeling epistemic states and their dynamics was presented by Spohn [58], although Spohn was more concerned with iterations of revision operations than with the Ramsey test and the triviality result. In this work, an epistemic state is not represented directly by a logical theory with an associated entrenchment, but indirectly, by an Ordinal Conditional Function (OCF), which defines an ordering on the models of the language by assigning an ordinal number to each model. Spohn uses only one operation, OCF revision, to represent all changes to epistemic states: expansion, contraction, revision and change of entrenchment. The result of such an operation is another OCF, allowing successive revisions can be performed. Spohn's approach is model-theoretic in that he provides no way of computing epistemic change operations: in Wobcke [64], we propose a syntactic analogue of Spohn's OCF revision and present algorithms for its computation.

In an AGM context in which belief states are represented by a theory with an epistemic entrenchment, various specific constructions for iterated theory change have been proposed, e.g., by Boutilier [4] and Nayak, Foo, Pagnucco and Sattar [50].

An interesting generalization is presented by Nayak [49] in which the new information for a revision is not a formula but a partial epistemic entrenchment. The problem with all these proposals is that the postulates are too strong. For example, Boutilier [4] uses the heuristic that the new information is always the least entrenched (he calls this “natural” revision), yet Nayak, Foo, Pagnucco and Sattar [50] use the heuristic that the new information is always more entrenched than any current nontrivial belief (“unnatural” revision?). An alternative approach to constraining iterated change is given by Darwiche and Pearl [8], who present four very general postulates intended to supplement the AGM postulates (see also Rott [57]). But whereas the specific constructions constrain iterated change too much, Darwiche and Pearl’s postulates fail to constrain iterated change sufficiently, since if the class of epistemic states is taken to be the theories of the corresponding conditional logic of validity, the postulates for iterated change will have no impact on the triviality result.

But all this fails to take account of another moral to be drawn from Spohn’s work: that there is no way to formally determine the strength of the new belief to be accepted under a revision. Intuitively, the same information could be obtained by different means, so the strength with which the new belief is held will not be a function of the belief’s content alone, but crucially dependent on the means of its acquisition. Thus one should not expect very strong postulates to constrain the process of iterated change if those postulates are to account for all possible ways in which the new belief can be acquired. In Spohn’s account, this intuition is captured by requiring that the input to a revision operation be a proposition plus an ordinal number representing the strength of the new information.

Here, we follow Spohn’s approach both in defining theory change operations on epistemic states and in assuming the new information comes with a given strength represented by an ordinal number. However, we adapt the AGM approach in representing epistemic states as entrenchments and in specifying postulates for belief change operations that capture a notion of minimal change of entrenchment. We provide a characterization of our theory change operations in terms of systems of spheres. Both for simplicity in presenting the construction, and because of its relevance to our computer implementation of an AGM belief revision system, we restrict attention to epistemic entrenchments that are ranked, as in Dixon and Wobcke [12] and Williams [62].

Definition 3.1 A *ranked epistemic entrenchment* is a (logically closed) belief set K together with a ranking function, assigning to each nontheorem a natural number known as its *rank*, that satisfies the following conditions.

- (RE1) If $A \vdash B$ then $rank(A) \leq rank(B)$.
- (RE2) $rank(A \wedge B) = \min(rank(A), rank(B))$.
- (RE3) $rank(A) = 0$ iff $A \notin K$.

A ranked epistemic entrenchment is a special kind of entrenchment defined by $A \leq B$ iff $rank(A) \leq rank(B)$ or $\vdash B$. Obviously, (RE1), (RE2) and (RE3) are the analogues of (EE2), (EE3) and (EE4); the other postulates follow from the properties of the natural numbers. The main effect of ranking an entrenchment is to rule out infinite descending chains of formulas that are less and less entrenched, and is thus related to what Lewis [40] calls the limit assumption.

Iterative AGM theory change operations on epistemic states can now be defined as follows. The idea is to interpret the principle of minimal change, formulated in Gärdenfors [21] as applying to the contents of a belief set, as also applying to the entrenchment of formulas in the belief set. This is justified by interpreting the rank of a belief as representing the strength with which a belief is held. According to this interpretation, if a belief is not affected by a belief change operation, its rank should not be affected either. That is, a belief which remains in the new belief set should retain its rank unless an inconsistency with the postulates of entrenchment forces it to be changed.

When defining an expansion or revision by a formula A , the intended new rank α of A must be given. So that such an operation is successful, we consider only expansions and revisions in which α is positive. The following special case arises: suppose A is ranked at β in K ; what should the rank of A be after expanding or revising to accept A with rank α where $\alpha < \beta$? One intuition, based on taking the information that A is of rank α to be definitive (Spohn's view), says that the rank should be α in the new state; another intuition, based on taking the new information as merely further confirmation of what is already known, says that A should remain ranked at β . In this paper, we adopt the second intuition: new information is incremental and so for an expansion or revision, the rank of a belief that remains in the revised state can only increase.¹

It is straightforward to make precise our notion of minimal change of entrenchment. Let K' be the state resulting from an expansion or revision of K (with ranking function $rank$) to accept A with rank α . The set of beliefs in K' is K_A^+ or K_A^* according to the AGM definitions. We define the ranking $rank'$ on the beliefs in K' as follows.

$$rank'(B) = \max(rank(B), \min(rank(A \rightarrow B), \alpha)) \text{ if } B \in K'.$$

Intuitively, the information A (at rank α) only provides *new* evidence for B if it can be combined with the evidence for $A \rightarrow B$ resulting in a new justification for B of greater strength than all prior justifications for B . The following consequence of this definition will simplify subsequent discussion: the ranks of formulas of the form $A \rightarrow B$ retained in K' are not affected by the operation, but the ranks of formulas of the form $A \vee B$ either stay the same or are increased to α .

If K_A^- is the contraction of K by A , it is even simpler to capture our notion of minimal change of entrenchment.

$$rank'(B) = rank(B) \text{ if } B \in K_A^-.$$

Another motivation for minimal change of entrenchment derives from the following postulates for expansion, contraction and revision, which are generalizations of the AGM postulates capturing a notion of “conservatively” accepting the new information. The key to expressing these postulates is to define a way to compare belief states in terms of the “amount” of information they contain. For AGM belief sets, this is easy: set-theoretic containment is the comparison measure. For ranked epistemic entrenchments, the following definition serves our purpose.

Definition 3.2 A ranked epistemic entrenchment K_1 (with ranking function $rank_1$) is *at least as conservative* as a ranked epistemic entrenchment K_2 (with ranking

function $rank_2$), denoted $K_1 \leq K_2$, if $K_1 \subseteq K_2$ and for all $A \in K_1$ such that $\not\vdash A$, $rank_1(A) \leq rank_2(A)$.

Note that this conservativeness ordering is compatible with set theoretic containment in that $K_1 \subseteq K_2$ is a necessary condition on $K_1 \leq K_2$. The conservativeness ordering is clearly a partial order. However, we need additional structure to state some of the postulates. Given two ranked epistemic entrenchments K_1 and K_2 , define $K_1 \wedge K_2$ (K_1 “meet” K_2 , generalizing set-theoretic intersection) as the set of beliefs in $K_1 \cap K_2$ with ranking function given by $rank(A) = \min(rank_1(A), rank_2(A))$, where $rank_1$ and $rank_2$ are the ranking functions of K_1 and K_2 , respectively. Given this definition, we could, alternatively, have defined $K_1 \leq K_2$ iff $K_1 \wedge K_2 = K_1$. The most important property of \leq is the descending chain condition which is inherited from the natural numbers: any sequence $K_1 \geq K_2 \geq \dots$ has a unique minimal element, i.e., an element K such that $K_i \geq K$ for all i . Equivalently, any subclass of ranked entrenchments which is closed under the meet operation has a unique minimal element. The maximal elements in the conservativeness ordering are those whose belief sets are inconsistent, and we shall not distinguish among these states.

We now specify six postulates for an expansion operation on epistemic states. In the following, $K_{A,\alpha}^+$ is the epistemic state resulting from K by expanding A to have rank α .

- (**K⁺1**) $K_{A,\alpha}^+$ is a belief state.
- (**K⁺2**) $\{A : \alpha\} \leq K_{A,\alpha}^+$.
- (**K⁺3**) $K \leq K_{A,\alpha}^+$.
- (**K⁺4**) If $A \in K$ and $rank(A) \geq \alpha$ then $K_{A,\alpha}^+ \leq K$.
- (**K⁺5**) If $K \leq H$ then $K_{A,\alpha}^+ \leq H_{A,\alpha}^+$.
- (**K⁺6**) $K_{A,\alpha}^+$ is the most conservative belief state satisfying (**K⁺1**)–(**K⁺5**).

The notation $\{A : \alpha\}$ in (**K⁺2**) indicates the belief state having A and all its nontheorem consequences at rank α , i.e., (**K⁺2**) says that A is of rank at least α in $K_{A,\alpha}^+$.

These postulates have the consequence that the content of $K_{A,\alpha}^+$ is $Cn(K \cup \{A\})$, as in the AGM definitions, and formulas are ranked in $K_{A,\alpha}^+$ the same as in K except for the consequences of A , whose ranks are increased to α if they are less than α in K .

Theorem 3.3 *If $K \neq K_\perp$, $K_{A,\alpha}^+$ is the AGM belief set K_A^+ with ranking function $rank_{A,\alpha}^+(B) = \max(rank(B), \min(rank(A \rightarrow B), \alpha))$ if $B \in K_A^+$. Otherwise $K_{A,\alpha}^+ = K_\perp$.*

We can also define the following postulates for contraction. Of course, unlike an expansion, it is necessary to specify only the formula to be removed.

- (**K⁻1**) K_A^- is a belief state.
- (**K⁻2**) $K_A^- \leq K$.
- (**K⁻3**) If $A \notin K$ or $K = K_\perp$ then $K_A^- = K$.
- (**K⁻4**) If $\not\vdash A$ and $K \neq K_\perp$ then $A \notin K_A^-$.
- (**K⁻5**) If $A \in K$ and $rank(A) = \alpha$ then $K \leq (K_A^-)_{A,\alpha}^+$.
- (**K⁻6**) If $\vdash A \leftrightarrow B$ then $K_A^- = K_B^-$.
- (**K⁻7**) $K_A^- \wedge K_B^- \leq K_{A \wedge B}^-$.

- (**K⁻8**) If $A \notin K_{A \wedge B}^-$ then $K_{A \wedge B}^- \leq K_A^-$.
 (**K⁻9**) K_A^- is the most conservative belief state satisfying (**K⁻1**)–(**K⁻8**).

The contraction K_A^- of K by A is the set K_A^- as defined by the AGM operations, with the ranks of the remaining formulas unchanged.

Theorem 3.4 *If $K \neq K_\perp$, K_A^- is the AGM belief set K_A^- with ranking function $rank_A^-$ defined by $rank_A^-(B) = rank(B)$ if $B \in K_A^-$. Otherwise $K_A^- = K_\perp$.*

Finally, we specify the following postulates for revision.

- (**K*1**) $K_{A,\alpha}^*$ is a belief state.
 (**K*2**) $\{A : \alpha\} \leq K_{A,\alpha}^*$.
 (**K*3**) $K_{A,\alpha}^* \leq K_{A,\alpha}^+$.
 (**K*4**) If $\neg A \notin K$ or $K = K_\perp$ then $K_{A,\alpha}^+ \leq K_{A,\alpha}^*$.
 (**K*5**) $K_{A,\alpha}^* = K_\perp$ only if $\vdash \neg A$ or $K = K_\perp$.
 (**K*6**) If $\vdash A \leftrightarrow B$ then $K_{A,\alpha}^* = K_{B,\alpha}^*$.
 (**K*7**) $K_{A \wedge B,\alpha}^* \leq (K_{A,\alpha}^*)_{B,\alpha}^+$.
 (**K*8**) If $\neg B \notin K_{A,\alpha}^*$ then $(K_{A,\alpha}^*)_{B,\alpha}^+ \leq K_{A \wedge B,\alpha}^*$.
 (**K*9**) If $\neg A \in K$ and $rank(\neg A) = \beta$ then $K \leq (K \wedge K_{A,\alpha}^*)_{\neg A,\beta}^+$.
 (**K*10**) $K_{A,\alpha}^*$ is the most conservative belief state satisfying (**K*1**)–(**K*9**).

Postulate (**K*2**) says that A must have rank of at least α in $K_{A,\alpha}^*$. (**K*3**) says that $K_{A,\alpha}^*$ contains no more beliefs and no beliefs more highly entrenched than if K is expanded by A with rank α . (**K*4**) says that $K_{A,\alpha}^*$ must contain at least as much information as $K_{A,\alpha}^+$ (and as a consequence of this and (**K*3**), if A has rank greater than α in K , A will still have this higher rank in $K_{A,\alpha}^*$). In interpreting (**K*5**), we consider as indistinguishable any two belief states whose belief sets are inconsistent. The postulate (**K*9**) is an analogue of the recovery postulate for contraction.

The revision of K to accept A with rank α is the AGM belief set K_A^* with the ranks of formulas $\neg A \vee B$ unchanged and ranks of formulas $A \vee B$ increased (if necessary) to α . That is, $K_{A,\alpha}^*$ is the most conservative belief state agreeing in ranking function with K on the set $K \cap K_A^*$ and having A ranked α .

Theorem 3.5 *If $K \neq K_\perp$, $K_{A,\alpha}^*$ is the AGM belief set $K_{A,\alpha}^*$ with ranking function $rank_{A,\alpha}^*$ defined by $rank_{A,\alpha}^*(B) = \max(rank(B), \min(rank(A \rightarrow B), \alpha))$ if $B \in K_{A,\alpha}^*$. Otherwise $K_{A,\alpha}^* = K_\perp$.*

As distinct from the AGM approach, the contraction and revision postulates *are* constructive and agree with the definition given above in terms of rank. Moreover, the Levi identity (in the form $K_{A,\alpha}^* = (K_{\neg A}^-)_{A,\alpha}^+$) and the Harper identity (in the form $K_A^- = K \wedge K_{\neg A,\alpha}^*$) both hold for epistemic state dynamics. However, the recovery postulate for contraction no longer holds in general: if $A \in K$ then $(K_A^-)_{A,\alpha}^+$ will have the same content as K , but the rank of A in $(K_A^-)_{A,\alpha}^+$ will be α rather than the rank of A in K . Notice also that whereas, as has been observed by Makinson [42], recovery for contraction plays no role in defining AGM revision, recovery does play a role in defining our revision operation.

Finally in this section, we generalize the characterization of theory change operations developed by Grove [32] to capture the dynamics of epistemic states. To illustrate this, an example of our iterative AGM belief revision is shown in Figure 1.




Figure 1: AGM Revision of Epistemic States

The belief set $K = Cn(\{a, b, c\})$ is to be revised to accept $\neg a$ with rank 10. As we shall elaborate in Section 5, nontheorems in $Cn(\{b\})$ are ranked 100, formulas in $Cn(\{a, b\})$ but not in $Cn(\{b\})$ are ranked 50, and formulas in $Cn(\{a, b, c\})$ but not in $Cn(\{a, b\})$ are ranked 30. Complete theories are described in terms of the polarity of the proposition symbols a , b and c that they contain, where \bar{x} indicates that the theory contains $\neg x$.

Figure 1 also illustrates a characterization of iterative AGM revision derived from Grove's [32] modeling using systems of spheres of complete theories, again restricted to the case of ranked epistemic entrenchments. For each rank r , there is a sphere S_r , said to be of rank r , consisting of all the complete theories containing all the formulas ranked higher than r , so that the rank of a nontheorem A is also the rank of the smallest sphere containing a complete theory containing $\neg A$. The sphere S_0 is the set of all complete theories containing the belief set. Call such a system of spheres a *ranked system of spheres*. Grove showed that the belief set resulting from a revision to accept a consistent formula A is represented by the set of all complete theories containing A that are contained in the smallest sphere containing at least one complete theory containing A (this is S_r where r is the rank of $\neg A$: call it S_A). To capture our iterative AGM revision operation, when accepting A with rank α , the complete theories containing A in S_A are shifted to rank 0 as in Grove's construction, while in addition, all the complete theories containing $\neg A$ ranked less than α are moved up to rank α . No other complete theory is affected.

Theorem 3.6 *If rank is a function determining a ranked system of spheres, the operation on ranked systems of spheres of complete theories defined by $rank^*_{A,\alpha}$, as follows, characterizes revision operations on consistent belief states that satisfy $(\mathbf{K}^*1) - (\mathbf{K}^*10)$.*

$$rank^*_{A,\alpha}(T) = \begin{cases} 0 & \text{if } A \in T \text{ and } T \in S_A \\ rank(T) & \text{if } A \in T \text{ and } T \notin S_A \\ \max(rank(T), \alpha) & \text{if } \neg A \in T \end{cases}$$

Thus in addition to capturing minimal change of entrenchment, our generalized AGM revision operation captures a strong notion of minimal change in the system of

spheres in the sense that only those theories that have to change rank do change rank.

In Figure 1, the revision to accept $\neg a$ with rank 10 results in the complete theories containing $\neg a$ from S_{50} moving down to S_0 and all the complete theories containing a ranked less than 10 moving up to S_{10} (in this example, there is only one such theory, in S_0).

For completeness, we present the analogous result for contraction.

Theorem 3.7 *If rank is a function determining a ranked system of spheres, the operation on ranked systems of spheres of complete theories defined by rank_A^- , as follows, characterizes contraction operations on consistent belief states that satisfy $(\mathbf{K}^{-1})-(\mathbf{K}^{-9})$.*

$$\text{rank}_A^-(T) = \begin{cases} 0 & \text{if } \neg A \in T \text{ and } T \in S_{\neg A} \\ \text{rank}(T) & \text{otherwise.} \end{cases}$$

So contraction also captures a principle of minimal change in the system of spheres.

4 A conditional logic of belief revision We now consider the relationship of belief revision to conditional logic. The main idea is that the dynamics of epistemic states acts as the semantics for a restricted logic which includes formulas of the form $A \Rightarrow B$ where A denotes a base belief. Recall that revisions of epistemic states take two arguments: a formula A and a rank α . By analogy to standard possible worlds semantics for conditional logic, $A \Rightarrow B$ is valid if for every epistemic state K and for all α , B holds at the state $K_{A,\alpha}^*$. A conditional $A \Rightarrow B$ holds at an epistemic state if B is accepted in the state resulting from the acceptance of A with some anticipated strength α . That is, a belief in $A \Rightarrow B$ is interpreted as a commitment to accept B conditional on A , but a commitment which is further conditional on A being acquired with the strength α . To make our semantics precise, we need to refer to sequences of revision operations: our definitions are similar to those used by Katsuno and Satoh [36] and Boutilier [4] to define conditional logics corresponding to belief revision operations.

Definition 4.1 *A belief formula is a formula built from the base logic and the propositional and conditional connectives using the following formation rules: if A is a base formula and B is a belief formula then $A \Rightarrow B$ is a conditional formula; if B is a conditional formula then $\neg B$ is a conditional formula; base formulas and conditional formulas are belief formulas and if B_1 and B_2 are belief formulas then so are $B_1 \wedge B_2$, $B_1 \vee B_2$, $\neg B_1$ and $B_1 \rightarrow B_2$.*

Definition 4.2 *A revision sequence is a sequence $\langle K_1, K_2, \dots, K_n \rangle$ such that for each $i < n$, $K_{i+1} = (K_i)_{A,\alpha}^*$ for some A and α .*

Definition 4.3 *A belief revision model is a tree τ whose branches are revision sequences such that if the root of the tree is the state K , for every base formula A there is one distinguished subtree $^*(\tau, A)$ whose root is $K_{A,\alpha}^*$ for some α : say also that the rank of the revision, $\rho(^*(\tau, A))$, is α . The selection function * must satisfy the following conditions.*

1. If $\vdash A \leftrightarrow B$ then $^*(\tau, A) = ^*(\tau, B)$.
2. If $^*(\tau, \neg A) = K_\perp$ then $^*(\tau, A) = \tau$.
3. If $^*(\tau, A) \neq K_\perp$ and $A \vdash B$ then $^*(\tau, B) \neq K_\perp$ and $\rho(^*(\tau, A)) \leq \rho(^*(\tau, B))$.

In order to avoid conflict with (**K*5**), a belief revision model τ is taken to be over the base logic whose theorems are those base formulas A such that $*(\tau, \neg A) = K_{\perp}$.

Definition 4.4 A belief formula B is valid if B is satisfied in all belief revision models.

A conditional belief $A \Rightarrow B$ is a commitment to accept B upon learning A with some expected degree of strength α as determined by the selection function. This choice of α will make no difference to the base formulas believed at the resulting state, but will to the conditional beliefs holding at that state. This is why $*$ is defined as mapping trees and formulas to trees, rather than belief states and formulas to belief states. Conditions (1)–(3) place some constraints on the expected strength of new information. Condition (1) says that two equivalent items of information must be expected to the same degree and must result in the same tree of sequences (hence conditional beliefs). Condition (2) says that no change to either belief state or revision sequences is allowed when accepting a belief regarded as necessarily holding at that state. Condition (3) says that if A is expected to degree α then any consequence B of A must be expected to at least the same degree, which is reasonable since the information A includes the information B .

We now define the semantics for our belief revision language. For a base formula A , A holds at a belief state K just in case A is contained in the belief set of K . For conditional formulas, we adapt the Ramsey test: $A \Rightarrow B$ holds at a belief state K at the root of a tree τ of revision sequences if B holds at the distinguished state selected by $*$ from τ and A . The complex case is that of formulas that mix the conditional arrow and propositional connectives. The easiest way to define these truth conditions is to reduce such formulas to equivalent base formulas, which can then be used to define the truth value of the original formula.

Definition 4.5 A belief state K satisfies a base formula B if B is contained in the belief set of K .

Definition 4.6 A belief state K which is the root of a tree τ of revision sequences satisfies $A \Rightarrow B$ if $*(\tau, A)$ satisfies B . A belief state K satisfies a conditional formula $\neg B$ if K does not satisfy B .

Definition 4.7 Let K be the root state of a belief revision model τ . Assume that the truth values in K of all conditional formulas which are subformulas of a belief formula B are known. Then B reduces in K to the base formula B' if B' is obtained by replacing subformulas of the form $B_1 \Rightarrow B_2$ by *true* if $K \models B_1 \Rightarrow B_2$ and by *false* otherwise. The truth conditions for all formulas can now be given as follows.

$K \models_{\tau} A$ if A is contained in the belief set of K for A a base formula.

$K \models_{\tau} B$ if $K \models_{\tau} B'$ where B reduces in K to B' .

Definition 4.8 A belief revision model satisfies a belief formula B if B holds at the root of the tree of revision sequences comprising the model.

The reason for the non-truth-functionality of the above definition is to account properly for formulas such as $((A \Rightarrow B) \wedge C) \vee ((A \Rightarrow B) \wedge D)$, which we take to have the same truth condition as $(A \Rightarrow B) \wedge (C \vee D)$: this holds at a state in which $A \Rightarrow B$

holds and in which $C \vee D$ is a belief (assuming C and D are base formulas). In particular, when $D \equiv \neg C$ the formula holds at every state in which $A \Rightarrow B$ holds. Note also that our semantics is not two-valued: an epistemic state may satisfy neither a base formula nor its negation. However, conditional formulas are all either true or false at an epistemic state, in contrast to the definitions of Gärdenfors [19] which allow belief sets to be agnostic about such formulas. The intuition underlying our definition is that, given a belief revision model, there are determinate facts about which beliefs hold at which states in which sequences, and hence to which conditional beliefs an agent is committed.

An equivalent definition can be given using van Fraassen's [60] supervaluations for base formulas. The following definition is closely related to the semantics proposed by Katsuno and Satoh [36] and will help explain the technical differences between the two approaches.

Definition 4.9 Let K be the root state of a belief revision model τ . Then K satisfies A if T satisfies A for every complete theory T containing K .

Definition 4.10 Let K be the root state of a belief revision model τ . A complete theory T_K containing K satisfies a belief formula B under the following truth conditions.

$$\begin{aligned} T_K \models_{\tau} A \Rightarrow B & \text{ if } K' \models_{\nu} B \text{ where } K' \text{ is the root of } \nu = *(\tau, A). \\ T_K \models_{\tau} \neg A & \text{ if } T_K \not\models_{\tau} A. \\ T_K \models_{\tau} A \wedge B & \text{ if } T_K \models_{\tau} A \text{ and } T_K \models_{\tau} B. \\ T_K \models_{\tau} A \vee B & \text{ if } T_K \models_{\tau} A \text{ or } T_K \models_{\tau} B. \\ T_K \models_{\tau} A \rightarrow B & \text{ if } T_K \not\models_{\tau} A \text{ or } T_K \models_{\tau} B. \end{aligned}$$

For the purposes of this definition, note that two identical complete theories T may have different truth assignments for a formula depending on which belief state K is under consideration.

We now present a conditional logic corresponding to our semantics for belief revision based on epistemic state dynamics. The logic **BR** is generated from the following axiom schemata and inference rules, where unless otherwise stated, A , B and C stand for base formulas. The formula $\Box A$ is defined as $\neg A \Rightarrow A$ for a base formula A .

- (B1) All truth functional tautologies.
- (B2) $A \Rightarrow A$.
- (B3) $\Box B \rightarrow (A \Rightarrow B)$.
- (B4) $(A \Rightarrow B) \rightarrow (A \rightarrow B)$.
- (B5) $(A \Rightarrow B) \wedge (A \Rightarrow C) \rightarrow (A \Rightarrow (B \wedge C))$ for B and C belief formulas.
- (B6) $(A \Rightarrow C) \wedge (B \Rightarrow C) \rightarrow ((A \vee B) \Rightarrow C)$.
- (B7) $\neg(A \Rightarrow \neg B) \rightarrow ((A \Rightarrow C) \rightarrow ((A \wedge B) \Rightarrow C))$.
- (B8) $\neg(A \Rightarrow \neg B) \rightarrow ((A \Rightarrow (B \rightarrow C)) \rightarrow (A \Rightarrow (B \Rightarrow C)))$.
- (B9) $((A \wedge B) \Rightarrow C) \leftrightarrow (A \Rightarrow ((A \wedge B) \Rightarrow C))$.
- (B10) $\neg\Box A \wedge (\neg A \Rightarrow (A \Rightarrow \neg B)) \rightarrow (((A \wedge B) \Rightarrow C) \leftrightarrow (\neg A \Rightarrow ((A \wedge B) \Rightarrow C)))$.
- (B11) $((A \wedge B) \Rightarrow C) \wedge (A \Rightarrow ((\neg A \vee \neg C) \Rightarrow C)) \rightarrow ((A \wedge B) \Rightarrow ((\neg A \vee \neg B \vee \neg C) \Rightarrow C))$.
- (B12) $((A \wedge B) \Rightarrow C) \wedge (A \Rightarrow \neg((\neg A \vee \neg C) \Rightarrow A)) \rightarrow$

- $$((A \wedge B) \Rightarrow \neg((\neg A \vee \neg B \vee \neg C) \Rightarrow (A \wedge B))).$$
- (B13) $(A \Rightarrow (B \vee C)) \rightarrow ((A \Rightarrow B) \vee (A \Rightarrow C))$ for B or C conditional formulas.
- (B14) $(A \Rightarrow \Box B) \rightarrow ((A \Rightarrow C) \leftrightarrow (A \Rightarrow (B \Rightarrow C)))$ for C a belief formula.
- (MP) From A and $A \rightarrow B$ infer B for A and B belief formulas.
- (RCEA) If $\vdash A \leftrightarrow B$ infer $(A \Rightarrow C) \leftrightarrow (B \Rightarrow C)$ for C a belief formula.
- (RCM) If $\vdash B \rightarrow C$ infer $(A \Rightarrow B) \rightarrow (A \Rightarrow C)$ for B and C belief formulas.

When considering only nonnested occurrences of the conditional operator, this logic is the same as Gärdenfors's logic **C** except for the centering axiom (A6).

$$(A6) \quad (A \wedge B) \rightarrow (A \Rightarrow B).$$

The reason that (A6) is included in **C** but not in **BR** is the difference in the definition of validity: Gärdenfors [19] takes a formula as valid if its negation is not contained in any consistent belief set, whereas a **BR** formula is valid if it holds at all belief revision models. To show the invalidity of (A6), consider an epistemic state that contains no nontrivial base formulas and in which $A \Rightarrow B$ is false: (A6) is true at such a state iff $\neg A \vee \neg B$ is a belief at that state (this is the base formula (A6) reduces to), but this is impossible by assumption (that such a state exists follows from completeness). Our version of (A6) is the following axiom, which follows from (B8).

$$(B8') \quad ((A \Rightarrow (B \wedge C)) \rightarrow (A \Rightarrow (B \Rightarrow C)))$$

We claim that (B8') captures (K*4) more accurately than (A6) does. (B8') says that if $B \wedge C$ holds on a revision to accept A , then on revising the resulting state to accept B , C still holds. While this follows from (A6) using (RCM), the difference is that in **BR** the formula $A \Rightarrow (B \wedge C)$ is determinately true or false at an epistemic state, so when A is a theorem, (B8') can be read as "if B and C hold at an epistemic state, then C holds at the state resulting from a revision to accept B ". However, (A6) must be read as " $B \wedge C \wedge \neg(B \Rightarrow C)$ does not hold in any epistemic state," which is a much weaker statement.

It is also interesting to see why (A6) is a theorem of the logic of Katsuno and Satoh [36]. Essentially, Katsuno and Satoh define satisfaction for all formulas A using our supervaluation condition: with a total pre-order on a set of complete theories, K satisfies A if T satisfies A for every complete theory T containing K . However, considering only complete theories, (A6) is perfectly reasonable. Our counterexample is blocked because there is no complete theory that does not contain $\neg A \vee \neg B$ while making $A \Rightarrow B$ false: if a complete theory does not contain $\neg A \vee \neg B$, it contains $A \wedge B$ and hence satisfies $A \Rightarrow B$.

The axioms (B9) and (B10) capture minimal change of entrenchment. Their effect is to ensure that the relative entrenchment of all formulas of the form $A \rightarrow B$ and $A \rightarrow C$ is the same in K_A^* as in K , and that the relative entrenchment of all formulas of the form $A \vee B$ and $A \vee C$ ranked higher than A in a consistent set K_A^* is the same in K_A^* as in K : formulas of this form not ranked higher than A in K_A^* are all ranked at the new rank α of A . The axioms (B11) and (B12) capture condition (3) on the strength with which new information can be expected. Their effect is to ensure that if $A \vdash B$ then when $C \in K_A^*$, if $B < C$ ($B \leq C$) in K_B^* then $A < C$ ($A \leq C$) in K_A^* . Hence suitable ranks for the revisions by A and B can be chosen to fulfill the condition.

The price to pay for the invalidity of **(A6)** is that either strong completeness or the deduction theorem must be given up, for certainly $\{A, B\} \models A \Rightarrow B$, so strong completeness would imply that $\{A, B\} \vdash A \Rightarrow B$ and then the deduction theorem that **(A6)** were a theorem. The logic **BR** satisfies the deduction theorem and is weakly complete. In proving this, it is useful to refer to belief sets with the properties of faithfulness and boundedness.

Definition 4.11 A set of **BR** sentences Γ is *faithful* if whenever $\Box A \in \Gamma$ for a base formula A , $A \Rightarrow B \in \Gamma$ iff $B \in \Gamma$ for all belief formulas B .

Definition 4.12 A set of **BR** sentences Γ is *bounded to degree n* if the language of Γ is a finite set of proposition symbols and the only conditional formulas are nested to depth at most n .

Theorem 4.13 *For every faithful, consistent set Γ of **BR** sentences bounded to degree n , there is a belief revision model which satisfies all and only the formulas of Γ .*

Corollary 4.14 ***BR** is sound and complete with respect to the belief revision models.*

Corollary 4.15 *There is a nontrivial belief revision model.*

Finally, although our approach to iterated revision has been formalized using conditional logic, there are obvious connections to work on nonmonotonic consequence relations, e.g., Gabbay [18], Makinson [43], Kraus, Lehmann and Magidor [37], Gärdenfors [23], Makinson and Gärdenfors [44], Lehmann and Magidor [38], and Fariñas del Cerro, Herzig and Lang [16]. In such work, properties of the consequence operations bear an exact correspondence to nonnested conditional formulas. Considering that Makinson and Gärdenfors [44] provide a translation between the AGM postulates and properties of nonmonotonic consequence operations, it should come as no surprise that the conditional logic versions of those properties agree with **BR** over the nonnested formulas. This formal connection between nonmonotonic consequence operations and conditional logic is further investigated by Bell [3], Arlo-Costa and Shapiro [2] and Crocco and Lamarre [6]. The work of Fariñas del Cerro, Herzig and Lang is also of interest because it addresses the question of what inferences hold given an incomplete ordering. The proposed answer is to define $A \vdash B$ if B holds at every belief state that results from accepting A using an entrenchment compatible with the ordering, giving a consequence operation that does not satisfy rational monotony (axiom **(B7)** in **BR**). Below, we answer this question using the most conservative entrenchment compatible with the given ordering, so that our consequence operation does satisfy rational monotony.

5 Computational belief change From the earliest work in theory change, e.g., Alchourrón and Makinson (1982), it has been recognized that any practical system using (in a strong sense) the operations of belief contraction and revision would have to operate using bases for belief sets. This has led to a large body of work defining analogues of the AGM postulates for use with theory bases, including Hansson [33], [34], [35], Fuhrmann [17], Nebel [51], Williams [61] and Nayak [48]. Much of this work seems concerned with reinterpreting the motivation for the AGM approach in terms of theory bases, in particular the notions of minimal change and the postulate of

recovery. But the whole motivation for base operations is unclear: if, as is agreed, the AGM approach is an attempt to model an idealized rational agent, what can be gained by proposing an idealization that works with the imperfect model of bases rather than with theories?

Put more concretely, whereas the notion of minimal change of theories has some intuitive plausibility, the idea that a rational revision operation should be a minimal change on a base is unmotivated, especially if it implies that a rational operation can be crucially dependent on the representation (presumably a contingent matter) of the belief set. That a base operation should “mimic” a theory change operation on its closure has been called by Dalal [7] the principle of irrelevance of syntax. All of the above base operations fail this principle. But if one is prepared to accept that a theory base extends to its most conservative entrenchment, it is possible to define AGM operations on bases that do satisfy the principle of irrelevance of syntax, not on belief sets but on belief states. Analogous considerations apply to the recovery postulate, the difference being that this *is* controversial from the point of view of theory change. Again, the fact that some base operations do not satisfy recovery is not a general argument in favor of base operations: Nayak [48] defines a base operation that does satisfy recovery on its closure. The upshot is that the base-sensitivity of theory change operations is in need of further justification. In what follows, we are therefore committed to the principle of irrelevance of syntax in requiring that an operation on a base is always a reflection of an operation on the belief state it represents.

The theory of the dynamics of epistemic states that we have developed generalizes the AGM approach to theory change. So far, this approach is not computational because it is not, in general, possible to represent epistemic states directly: an epistemic state involves the assignment of a rank to each formula, and if there are infinitely many different ranks then the state cannot be finitely represented. A related concern is that incomplete information about epistemic states, including their ranking functions, is a fundamental aspect of problems in default reasoning. According to our approach, which uses the E-bases of Rott [56] and Williams [61], an E-base is taken to represent the most conservative entrenchment compatible with the initial ordering, as in Wobcke [63]. We define operations on E-bases that reflect the epistemic state change operations of expansion, contraction, and revision based on minimal change of entrenchment. These operations satisfy the principle of irrelevance of syntax in the sense that any two E-bases that have the same most conservative entrenchment will, when expanded, contracted, or revised, result in E-bases generating the same most conservative entrenchment. But obviously, expanding, contracting, or revising two different E-bases which have only the same contents will not necessarily result in E-bases with the same contents, so looking only at the contents of belief states, the principle of irrelevance of syntax will not be satisfied.

5.1 Bases for epistemic entrenchments Epistemic entrenchments themselves are not suitable for computation because too much information (the entrenchment relation between *every* pair of formulas) has to be provided. Rott [56] investigated the use of so-called *E-bases* for specifying entrenchments. The formulation we use is the special case of an E-base in which a belief state is represented by a theory base together with a function assigning to each element of the base a natural number known as its

rank, as considered by Williams [61]. The given ranking on the base formulas is understood to extend in a unique way to the most conservative entrenchment compatible with the ranking, (see Wobcke [63]). Recall that since the conservativeness ordering satisfies the descending chain condition, such a unique most conservative entrenchment always exists whenever the base is consistent. Note also that the nonredundancy condition on E-bases is not essential for theoretical purposes but is useful for computational efficiency.

Definition 5.1 A *ranked E-base* Γ is a set of formulas together with a ranking function, assigning to each nontheorem a natural number known as its *rank*, that satisfies the following (nonredundancy) condition.

(R) For all $A \in \Gamma$, $\{B \in \Gamma \mid \text{rank}(B) > \text{rank}(A)\} \not\vdash A$.

Definition 5.2 An epistemic entrenchment \leq (with strict part $<$) is *compatible* with an E-base Γ if for all A and B contained in Γ , if $\text{rank}(A) \leq \text{rank}(B)$ then $A \leq B$ and if $\text{rank}(A) < \text{rank}(B)$ then $A < B$.

In the present context, the use of most conservative entrenchments, which we motivated above in terms of minimal change of entrenchment, now embodies an assumption of foundationalism. The rank of a formula is intended to represent the degree of evidence possessed by the system for a particular belief. A representation of a belief state is therefore a collection of formulas (those of the base) together with the degrees of evidence for each formula in the base. The belief set represented is the logical closure of the belief set base. Similarly, the most conservative entrenchment is the “closure,” in some sense, of the E-base. To define this sense of closure, we take it that the only evidence possessed by the system for a belief is that derived from the evidence for formulas in the base: the evidence for a formula derived from a nonredundant collection of formulas is the evidence for the conjunction of these formulas, which is just the evidence for the least entrenched formula in the conjunction. In this sense, the generated entrenchment is conservative in not attributing evidence for a belief other than that warranted by the evidence for the base beliefs from which that belief derives.

Williams [61] presents the following characterization of the entrenchment generated from a ranked E-base. Let the E-base Γ be partitioned into subsets $\Gamma_1, \Gamma_2, \dots$, where Γ_r is the set of base formulas with rank r and let $\overline{\Gamma}_r = \bigcup_{i=r}^{\infty} \Gamma_i$. Define a ranked epistemic entrenchment by setting $\text{rank}(A) = \max(\{r \mid \overline{\Gamma}_r \vdash A\})$ when $A \in \text{Cn}(\Gamma)$ and $\text{rank}(A) = 0$ when $A \notin \text{Cn}(\Gamma)$.

Definition 5.3 An E-base Γ is a *base* for a ranked epistemic entrenchment K if K is the most conservative ranked epistemic entrenchment compatible with Γ .

Theorem 5.4 The E-base Γ is a base for a ranked epistemic entrenchment K iff K is defined by setting $\text{rank}(A) = \max(\{r \mid \overline{\Gamma}_r \vdash A\})$ when $A \in \text{Cn}(\Gamma)$ and $\text{rank}(A) = 0$ when $A \notin \text{Cn}(\Gamma)$.

It is an obvious corollary that in terms of Grove’s systems of spheres, the most conservative entrenchment compatible with an E-base is that represented by allowing each complete theory to “sink” to the lowest possible rank while still respecting ranks of the formulas in the base. This provides a strong connection to the rational closure operation of Lehmann and Magidor [38].

Definition 5.5 A ranked system of spheres *represents* an E-base Γ if for each non-theorem $A \in \Gamma$, $\text{rank}(A)$ is the rank of the smallest sphere containing a complete theory containing $\neg A$.

Corollary 5.6 *The most conservative ranked epistemic entrenchment compatible with an E-base Γ is represented by the ranked system of spheres defined by setting $\text{rank}(T)$ to be the smallest r such that T contains A for all $A \in \Gamma$ ranked greater than r .*

5.2 Algorithms for epistemic state change E-bases are ideal for the direct representation of epistemic states, and in this section, we show how they can be used as the basis of algorithms for the computation of the iterative theory change operations of expansion, contraction and revision. The algorithms are completely independent of the base logic; what is assumed is the existence of a theorem prover for the base logic which can be used by our algorithms. Our algorithms terminate (i.e., they are algorithms) provided the theorem prover is guaranteed to terminate on both consequences and nonconsequences of a theory. Assuming an oracle for the theorem prover, all the algorithms have run-time complexity linear in the size of the base. But in practice, the most expensive part of the algorithms is the use of the theorem prover; hence our algorithms are written to minimize the number of times the theorem prover is called. A particular implementation of our algorithms over the base logic of first order logic with equality (along with some heuristic methods for handling nontermination of the theorem prover) is described in [12].²

5.2.1 Determination of rank The computation of the rank of a formula A is the simplest process to define mathematically, but it is where the complexity of the system lies in practice. This is so because rank determination is far worse than theorem proving since not just one proof, but the “best” proof of A (the one with the highest rank) must be found. The following procedure uses a variant of the branch-and-bound algorithm. It simply finds proofs of A in turn and associates with each proof a rank defined as the rank of the lowest ranked formula involved in the proof. This is a lower bound for the rank of A . Subsequent attempts to prove A are terminated as soon as a formula of equal or lower rank is used. When there are no further proofs of A greater than the cutoff, the greatest lower bound of the ranks is returned as the rank of A .³ If no proof is found, a rank of 0 is returned, indicating that A is not a belief.

Suppose **proof**(Γ, r, A) denotes a procedure to find all proofs of the formula A using only the formulas from the base Γ of rank at least r , i.e., using only the formulas in $\overline{\Gamma}_r$. The procedure **rank**(Γ, A), written below in pseudo-code, repeatedly calls **proof** with an increasing rank bound to eventually return the least upper bound of all proofs of A .

```

rank( $\Gamma, A$ )
   $r := 0$ 
  while  $p := \mathbf{proof}(\Gamma, r + 1, A)$  do      % more proofs exist
     $r := \mathbf{rank}(p)$ 
  return( $r$ )

```

5.2.2 Expansion We implement the expansion procedure $\mathbf{expand}(\Gamma, A, \mathit{newrank})$ using the following algorithm, where Γ is the base of the belief set, A is the belief being added to the base, and $\mathit{newrank}$ (assumed positive) is the intended rank of A . The procedure $\mathbf{update}(\Gamma, A, \mathit{newrank})$ either adds the formula A to the base Γ with rank $\mathit{newrank}$ if A is not already contained in the base or else changes the rank of A to $\mathit{newrank}$, $\mathbf{delete}(\Gamma, A)$ deletes A from the base Γ , and $\mathbf{theorem}(A)$ uses the theorem prover to test whether A is a theorem.

```

expand( $\Gamma, A, \mathit{newrank}$ )
  if theorem( $A$ ) or (rank( $\Gamma, \neg A$ ) > 0)
    return( $\Gamma$ )
  else
     $\mathit{oldrank} := \mathbf{rank}(\Gamma, A)$ 
    if ( $\mathit{newrank} \leq \mathit{oldrank}$ )
      return( $\Gamma$ )
    else
       $\Delta := \mathbf{update}(\Gamma, A, \mathit{newrank})$ 
      for each  $B \in \Gamma$  of rank  $r$  such that  $\mathit{oldrank} \leq r \leq \mathit{newrank}$  do
        if proof( $\Delta - \{B\}, r, B$ )    %  $B$  is redundant
           $\Delta := \mathbf{delete}(\Delta, B)$ 
      return( $\Delta$ )

```

Some comments on this algorithm are in order. First, when the base is to be expanded by a formula A , it is tested for logical consistency with the current base and the expansion rejected if it is not consistent. This is in conflict with our earlier definitions, but this is justified on the basis that inconsistent belief sets, especially over classical logics, are not useful in practice. Second, if the formula is consistent with the belief set, its rank is determined for consistency with the rest of the generated entrenchment. The given new rank $\mathit{newrank}$ is taken to indicate further evidence for A , so if A already has a rank of at least $\mathit{newrank}$, no change to Γ is made. If the rank of A in Γ is less than $\mathit{newrank}$, A is explicitly added to the base with rank $\mathit{newrank}$ and any necessary changes to the base are made. The changes are to ensure that the new base contains no redundancy. Any formula in Γ with rank at least the old rank of A but not greater than $\mathit{newrank}$ may become redundant by the change in rank of A , in particular if it is derivable with the same or higher rank from the other formulas in the new base after A has been added. In this case, it is deleted. Note that determining the rank of an element in the base does not require the prover but is merely a lookup of its stored rank in the old base Γ . Note also that for each iteration of the loop, the test for whether a base formula B is derivable uses the base which may have been reduced after previous iterations of the loop have deleted some formulas. This does not affect the end result because a formula C is deleted only if it is derivable from the formulas in the new base of at least the same rank, so B is derivable from the new base with some rank if it is also derivable from the new base with the same rank after C has been deleted. This means that the expansion algorithm does not need a copy of the original base Γ but instead can repeatedly update the one database.

The following theorem guarantees that our expansion algorithm correctly computes an E-base for the belief state $K_{A,\alpha}^+$ as defined using the postulates (\mathbf{K}^+1) – (\mathbf{K}^+6) .

Theorem 5.7 *If Γ is a base for $K \neq K_{\perp}$, $\mathbf{expand}(\Gamma, A, \alpha)$ is a base for $K_{A,\alpha}^+$.*

5.2.3 Contraction We now present an algorithm $\mathbf{contract}(\Gamma, A)$ for contraction of the base Γ by the formula A .

```

contract( $\Gamma, A$ )
  if theorem( $A$ )
    return( $\Gamma$ )
  else
     $\Delta := \Gamma$ 
     $oldrank := \mathbf{rank}(\Gamma, A)$ 
    for each  $B \in \Gamma$  of rank  $r$  such that  $r \leq oldrank$  do
      if not(proof( $\Delta, oldrank + 1, A \vee B$ ))    % ( $C^-$ ) condition
         $\Delta := \mathbf{delete}(\Delta, B)$ 
        if ( $r < oldrank$ ) or not(proof( $\Delta, oldrank + 1, A \rightarrow B$ ))
           $\Delta := \mathbf{update}(\Delta, A \rightarrow B, r)$ 
    return( $\Delta$ )

```

If an attempt is made to contract a theorem, the attempt is rejected in accordance with the AGM definitions. A change is made only if A is derivable from the base. In this case, each member B of the base is checked to see whether B is deleted from the base, and further, if $A \rightarrow B$ is to be added to the base. By the (C^-) condition, B is removed from the base if $rank(A) = rank(A \vee B)$. The formula $A \rightarrow B$ needs to be added to the base if it will not be derivable in the new base: its rank is that of B in the original base. The test for whether $A \rightarrow B$ is derivable in the new base can be greatly simplified. In the case where B is ranked less than A , given that $rank(A) = rank(A \vee B)$, $A \rightarrow B$ must be ranked the same as B by (RE2), hence $A \rightarrow B$ cannot be derivable after B is deleted since the original base Γ is not redundant. The only other case is where A , B and $A \vee B$ are all ranked the same in the original state. In this case, $A \rightarrow B$ must be added to the base with that rank if it is not already derivable from the set of higher ranked formulas. Note that this algorithm repeatedly tests the ranks of formulas in the constructed base Δ rather than in the original base Γ as the (C^-) condition requires. Again, as with expansion, this does not affect the result because of the nature of the tests performed. The first tests whether there is a proof of $A \vee B$ using only formulas ranked higher than A ; the second tests whether there is a proof of $A \rightarrow B$ using the same set of formulas. Since none of these formulas are affected by earlier iterations of the loop, this set of formulas is the same in each Δ as in the original Γ .

The replacement of B by $A \rightarrow B$ is necessary because of the AGM postulate ($\mathbf{K}^- \mathbf{5}$),⁴ $K \subseteq (K_A^-)_A^+$, from which we can derive that if $B \in K$ then $K_A^- \vdash A \rightarrow B$ (if $B \in K$, $K_A^- \cup \{A\} \vdash B$ by ($\mathbf{K}^- \mathbf{5}$), so $K_A^- \vdash A \rightarrow B$ by the deduction theorem). This replacement is redundant if $A \rightarrow B \in K$. The rank for the added formula $A \rightarrow B$ is that of B in the original base Γ because, given that $A \rightarrow B$ is not provable from the set of formulas ranked higher than B , its rank in Γ is the same as that of B in Γ . Note that the contraction of A is successful even if A is not contained in the base because at least one base formula from each proof of A is removed.

The following theorem guarantees that our contraction algorithm correctly computes an E-base for the belief state K_A^- as defined using the postulates ($\mathbf{K}^- \mathbf{1}$) – ($\mathbf{K}^- \mathbf{9}$).

Theorem 5.8 *If Γ is a base for $K \neq K_{\perp}$, $\mathbf{contract}(\Gamma, A)$ is a base for K_A^- .*

5.2.4 Revision Finally, we present an algorithm $\mathbf{revise}(\Gamma, A, \mathit{newrank})$ for computing the revision of an epistemic state Γ to accept A at a positive rank $\mathit{newrank}$.

$\mathbf{revise}(\Gamma, A, \mathit{newrank})$

$\mathit{oldrank}_A := \mathbf{rank}(\Gamma, A)$

if $\mathbf{theorem}(A)$ or $\mathbf{theorem}(\neg A)$ or $(\mathit{newrank} \leq \mathit{oldrank}_A)$

 return(Γ)

else

$\Delta := \Gamma$

$\mathit{oldrank}_{\neg A} := \mathbf{rank}(\Gamma, \neg A)$

 for each $B \in \Gamma$ of rank r such that $r \leq \mathit{oldrank}_{\neg A}$ do

 if $(r \leq \mathit{newrank})$ or not($\mathbf{proof}(\Delta, \mathit{oldrank}_{\neg A} + 1, A \rightarrow B)$) % (C*) cond.

$\Delta := \mathbf{delete}(\Delta, B)$

 if $(r > \mathit{newrank})$ and $((r < \mathit{oldrank}_{\neg A})$ or not($\mathbf{proof}(\Delta, r + 1, A \vee B)$))

$\Delta := \mathbf{update}(\Delta, A \vee B, r)$

$\Delta := \mathbf{expand}(\Delta, A, \mathit{newrank})$

 return(Δ)

If A is a theorem or a contradiction, no change is made to the base. Otherwise any formula B such that $\neg A$ is of rank equal to $A \rightarrow B$ in the original base is replaced by $A \vee B$ where this is not redundant. It suffices to consider only those B ranked less than or equal to $\neg A$ because if $\neg A < B$ then $\neg A < A \rightarrow B$, so B remains in the revised set. The replacement formula is redundant if its rank is at most the new rank of A (because $A \vee B$ is a consequence of A) or if $A \vee B$ can be proven from the base with a higher rank. The final step in the algorithm is to expand the remaining base by A . Note, again, that the repeated calls to \mathbf{proof} use only the formulas ranked higher than $\neg A$, and this set does not change with repetitions of the loop.

The following theorem guarantees that our revision algorithm correctly computes an E-base for the belief state $K_{A,\alpha}^*$ as defined using the postulates (\mathbf{K}^*1)–(\mathbf{K}^*10).

Theorem 5.9 *If Γ is a base for $K \neq K_{\perp}$, $\mathbf{revise}(\Gamma, A, \alpha)$ is a base for $K_{A,\alpha}^*$.*

Finally, we briefly discuss the complexity of the revision algorithm in terms of its use of the theorem prover. The number of calls to the procedure \mathbf{proof} depends on the values of $\mathit{oldrank}$ (the initial rank of $\neg A$) and $\mathit{newrank}$. Suppose N is the size of the base, N_r the number of base formulas ranked at most r and n_r the number ranked exactly r . Then, considering only those cases where A is not initially contained in the base, in the worst case, AGM revision requires $N_{\mathit{newrank}} - N_{\mathit{oldrank}} + n_{\mathit{oldrank}}$ calls to \mathbf{proof} if $\mathit{newrank} \geq \mathit{oldrank}$ and $N_{\mathit{oldrank}} - N_{\mathit{newrank}}$ calls to \mathbf{proof} if $\mathit{newrank} < \mathit{oldrank}$. Thus the complexity of iterative AGM revision in terms of the use of the theorem prover is proportional to the degree of change made to the base, and this means that not only is iterative AGM revision a “minimal” change in the theoretical sense, but that it can be implemented using a “minimal” change in the computational sense. The most significant fact from a practical point of view is that the size of the base increases by at most one, which means that the base grows linearly with the number of iterations. With the implementation of Spohn’s [58] OCF revision given in Wobcke [64],

the size of the base may double at each iteration, making iterated OCF revisions much more expensive than iterated AGM revisions.

6 Nonmonotonic reasoning We now investigate the relationship of nonmonotonic reasoning to both belief revision and conditional logic. The basic idea behind approaches to nonmonotonic reasoning based on belief revision, e.g., Gärdenfors [23], Wobcke [63], is that a default represents a policy of belief change. More precisely, a default such as $bird(tweety) \Rightarrow fly(tweety)$ represents a commitment to accept the belief $fly(tweety)$ upon learning that $bird(tweety)$. Such a commitment is understood as being with respect to an epistemic state that corresponds to a given initial default theory, which is usually fixed for the purpose of analysis. Note that if the initial collection of beliefs is to be consistent, there can be no “universal” defaults such as $\forall x(bird(x) \rightarrow fly(x))$ if in addition, some instance of the rule contradicts the universal statement, i.e., some particular bird is known not to fly. Thus under the belief revision approach to nonmonotonic reasoning, a generic default such as “birds fly” is represented as the formula $\forall x(bird(x) \rightarrow fly(x))$, which represents a collection of belief revision policies $bird(t) \Rightarrow fly(t)$, one for each term t in the language. For any exceptional bird such as the nonflying *tweety*, the universal default must be weakened to $\forall x((x \neq tweety) \wedge bird(x) \rightarrow fly(x))$ to ensure the consistency of the belief set.

Belief revision has already been shown to have a close intuitive connection to nonmonotonic reasoning in Gärdenfors [22]. The connection between belief revision and nonmonotonic consequence operations was explored in Makinson and Gärdenfors [44] and extended by Gärdenfors [23] who introduced the concept of an expectation (the analogue of an entrenchment) and showed how belief revision relates to the default reasoning system of Poole [54]. All these connections are based on the observation that belief revision is nonmonotonic, i.e., it is possible that $B \in K$ while $B \notin K_A^*$. Thus if a consequence operation \sim is defined by $A \sim B$ iff $B \in K_A^*$, then \sim is nonmonotonic. However, there is more at issue in nonmonotonic reasoning than just nonmonotonicity of the consequence operation. We claim that a central feature of the core problems of interest is the partiality or incompleteness of the initial information. So expectations cannot be regarded as a solution to the problems of nonmonotonic reasoning because too much information has to be given initially, e.g., a total pre-order on all formulas of the language.

The insufficiency of expectations alone to address problems in nonmonotonic reasoning can be seen by examining a simple problem in inheritance. Consider the theory: birds typically fly, robins and penguins are birds, penguins typically don't fly, Tweety is a bird and Opus is a penguin. Some desired conclusions are that robins typically fly, that Tweety can fly but Opus cannot fly and that birds colored black can fly. What can be shown is that *some* expectation that gives these conclusions can be defined, but not how this expectation is to be generated from just the initial default theory. Partial or incomplete orderings on beliefs must be considered.

Incomplete orderings on formulas in a logical language have been proposed independently as a basis for implementing nonmonotonic reasoning systems. These approaches, stemming from the intuitions of probability theory, bear a close relationship to those using entrenchments to represent belief revision operations. System Z, in Pearl [52], uses a total pre-order on a finite collection of defaults to define ϵ -

entailment, the “core” of nonmonotonic inference, and 1-entailment, an extension of ϵ -entailment which has close connections to the rational nonmonotonic consequence operations of Lehmann and Magidor [38], as shown by Goldszmidt and Pearl [30]. A number of perceived weaknesses of 1-entailment have been addressed in subsequent work. Goldszmidt, Morris and Pearl [29] define an approach based on “maximum entropy” orderings of defaults, while Geffner and Pearl [27] define conditional entailment, which allows partial orderings of defaults. In these approaches, the ordering of the defaults can be generated *automatically* from a given collection of defaults. This represents a major difference between approaches based on orderings and those based on belief revision because the entrenchment or expectation used by a belief revision system is assumed to be given.

The computational approach to the dynamics of epistemic states that we have developed allows the possibility of incompleteness in the specification of default theories: a default theory is taken to stand for the most conservative ranked epistemic entrenchment that respects the defaults. However, this interpretation works only for propositional default theories, and, moreover, our system corresponds to the 1-entailment of Pearl [52] and the rational consequence of Lehmann and Magidor [38], and thus yields the same counterintuitive results. Concluding that additional information is needed to avoid these results, we require a first order default theory to be defined by a ranked E-base, analogous to the “variable-strength” defaults of Goldszmidt and Pearl [31]. A ranked E-base is taken as standing for the most conservative ranked epistemic entrenchment compatible with the base, giving a nonmonotonic consequence operation that satisfies rational monotony, which corresponds to **(K*8)**. Rational monotony is central in the AGM account of belief revision, although work by Rott [57] shows how a generalized form of epistemic entrenchment can be used as a representation for a class of revision operations that does not, in general, satisfy rational monotony. Our approach to incomplete information is closely related to that described in Fariñas del Cerro, Herzig and Lang [16], who take a given base as extending to all expectations that are compatible with it, also giving a nonmonotonic consequence operation that does not satisfy rational monotony.

To further motivate our approach to nonmonotonic reasoning, we first consider the connections to conditional logic. In nonmonotonic reasoning, what is considered are the *extensions* of a default theory, i.e., sets of sentences containing the theory that include as many default conclusions as are consistently possible. In section 6.1, we show a connection between extensions of a default theory expressed in conditional logic and the most conservative epistemic state representing the theory. This connection is limited to propositional default theories. We propose to handle first order default theories by adopting the additional assumptions of uniqueness of names and independence of default instances, as described in section 6.2. We present an algorithm for the computation of belief revision operations which incorporates the independence of default instances, whereas uniqueness of names is assumed to be implemented as part of the theorem prover, as it has been in the system described in Dixon and Wobcke [13]. In section 6.3, we show the behavior of the system on a number of benchmark problems in nonmonotonic reasoning collected by Lifschitz [41].

6.1 Nonmonotonic reasoning and conditional logic The simplest approach to

nonmonotonic reasoning using conditional logic is to use deduction in conditional logic to derive conclusions nonmonotonically. A default such as that birds fly is represented as a *collection* of beliefs $bird(x) \Rightarrow fly(x)$, one such belief for each x . Using the Ramsey test for conditionals, this is taken to be a set of belief revision policies, i.e., “if I come to accept that x is a bird, I will also accept that x flies.” This indicates that we would conclude “by default” $fly(x)$ from $bird(x)$. Thus given a set of beliefs Γ , the set of nonmonotonic conclusions derivable from a fact A is the set the beliefs that would result were A to be accepted as a belief, which is just the set K_A^* as defined by the AGM postulates. This is also the set of B such that $A \Rightarrow B$ is contained in the default theory Γ .

Unfortunately, this by itself is insufficient to generate all the desired conclusions from an initial set of defaults. As Delgrande [10] argues, the crux of the problem is representing the independence or irrelevance of various facts (in relation to particular defaults). For example, if we are to conclude the default that birds that have wings fly based on the default that birds fly, we need to know that a bird’s having wings does not affect its ability to fly. Delgrande achieves this by using specially defined procedures for constructing extensions of default theories. We develop an approach to nonmonotonic reasoning in which assumptions of independence or irrelevance are explicitly expressed in conditional logic, with extensions defined by logical consequence. This approach has close connections to Pearl’s [52] System Z and to Lehmann and Magidor’s [38] rational closure of a conditional knowledge base.

Definition 6.1 A *default theory* is a finite set of **BR** formulas of the form $A \Rightarrow B$ where A and B are base formulas.

Definition 6.2 A formula B is *independent* of A in a theory Γ if $\neg(A \Rightarrow \neg B) \in \Gamma$.

Intuitively, if I come to believe B , this has no effect on my beliefs that follow from A by default in the sense that I do not have to revise any of these beliefs in order to accept A . The importance of independence assumptions in nonmonotonic reasoning is that they sanction both strengthening the antecedent and transitivity of defaults, which enable default conclusions to be propagated. More precisely, the following are theorems of the logic **BR** (here A , B and C are base formulas).

$$\begin{aligned} \neg(A \Rightarrow \neg B) &\rightarrow ((A \Rightarrow C) \rightarrow ((A \wedge B) \Rightarrow C)) \\ \neg(B \Rightarrow \neg A) &\rightarrow ((A \Rightarrow B) \wedge (B \Rightarrow C) \rightarrow (A \Rightarrow C)) \end{aligned}$$

These formulas are consequences of our axiom **(B7)**, corresponding to the AGM postulate **(K*8)**, which also arises out of a need to capture the independence of beliefs.

Definition 6.3 An assumption $\neg(A \Rightarrow \neg B)$ is *at least as strong as* $\neg(A' \Rightarrow \neg B')$ if $A' \vdash A$.

Intuitively, the stronger the independence assumption, the “more” cases of strengthening the antecedent and transitivity it allows, in the sense that whenever an object has the property A' it has the property A , so the class of A -objects includes at least all the A' -objects.

Definition 6.4 A set of independence assumptions S is stronger than a set S' if S contains one or more equally strong independence assumptions A not contained in S' and for any other independence assumption A' at least as strong as A , $A' \in S$ iff $A' \in S'$.

Definition 6.5 An *extension* of a BR theory Γ is a faithful, consistent set of BR sentences Γ^* containing Γ such that no such set includes a stronger set of independence assumptions than Γ^* .

Since consistent sets of BR sentences are not, in general, faithful, they do not correspond to entrenchments. However, a consistent set of BR sentences Γ can be associated with a number of entrenchments, one for each faithful, consistent set containing Γ , in which $A \leq B$ for base formulas A and B iff $\Box B \vee \neg((\neg A \vee \neg B) \Rightarrow A)$ is contained in Γ . Note that the formula $\neg((\neg A \vee \neg B) \Rightarrow A)$ is an independence assumption, specifically, one that represents the independence of $\neg A$ from $\neg A \vee \neg B$. We now show that the extensions of a theory Γ correspond to the most conservative epistemic states that agree with the defaults in Γ , and hence that a consistent default theory has exactly one extension. This is analogous to results presented by Pearl [52] for System Z and Lehmann and Magido [38] for rational closure (the correspondence between these two systems was shown by Goldszmidt and Pearl [30]).

Definition 6.6 A default theory Γ *represents* an entrenchment relation $A \leq B$ (respectively $A < B$) if Γ contains $\Box B \vee \neg((\neg A \vee \neg B) \Rightarrow A)$ (respectively $\neg \Box A \wedge ((\neg A \vee \neg B) \Rightarrow B)$).

Definition 6.7 A ranked epistemic entrenchment *respects* a BR theory Γ if it satisfies all entrenchment relations (over \leq and $<$) represented in Γ .

Theorem 6.8 *Extensions of a BR theory Γ are in correspondence with the most conservative ranked epistemic entrenchments respecting Γ .*

Corollary 6.9 *A consistent default theory has a unique extension.*

Thus the formation of an extension of a default theory has a direct correlate in epistemic entrenchment. The rule is: given a default theory, assume that the beliefs are ranked as low as possible consistent with the defaults. The default conclusions B that follow from a formula A with respect to the default theory Γ are those formulas B such that $A \Rightarrow B$ is contained in every extension of Γ , or equivalently, those formulas B that hold at those states which result from a revision to accept A of every belief state K that is a most conservative belief state respecting Γ .

Definition 6.10 $A \vdash_{\Gamma} B$ if $A \Rightarrow B$ is contained in every extension of Γ .

To illustrate the use of independence assumptions, or equivalently, most conservative entrenchments, in generating nonmonotonic conclusions, consider the simple default theory with the law $\Box(\text{penguin}(x) \rightarrow \text{bird}(x))$ and the two defaults $\text{bird}(x) \Rightarrow \text{fly}(x)$ and $\text{penguin}(x) \Rightarrow \neg \text{fly}(x)$. By (B4), our default theory contains the formulas $\text{bird}(x) \rightarrow \text{fly}(x)$ and $\text{penguin}(x) \rightarrow \neg \text{fly}(x)$, and therefore contains the formula $\neg \text{penguin}(x)$: i.e., anything is, by default, not a penguin. An E-base for the belief state corresponding to this simple theory is shown in Figure 2. The formula $\text{penguin}(x) \rightarrow \text{bird}(x)$ is treated as a theorem of the base logic and is ranked at the highest level. The defaults $\text{bird}(x) \Rightarrow \text{fly}(x)$ and $\text{penguin}(x) \Rightarrow \neg \text{fly}(x)$ are represented by the entrenchments $\neg \text{bird}(x) < \text{bird}(x) \rightarrow \text{fly}(x)$ and $\neg \text{penguin}(x) < \text{penguin}(x) \rightarrow \neg \text{fly}(x)$: note that $\neg \text{bird}(x)$ is a nonbelief and hence is less entrenched than any belief and that $\neg \text{penguin}(x)$ is ranked at the level of $\text{bird}(x) \rightarrow \text{fly}(x)$ because it follows from that formula together with $\text{penguin}(x) \rightarrow \neg \text{fly}(x)$ using the theorem $\text{penguin}(x) \rightarrow \text{bird}(x)$. Because $\neg \text{penguin}(x)$ is a belief, there must be at least

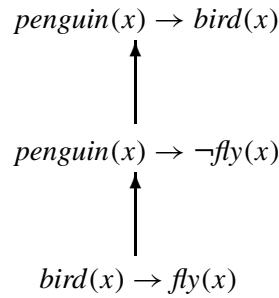


Figure 2: E-base for Simple Default Theory

two levels of ranking to generate both desired conclusions, and since there are exactly two levels of ranking of the nontheorems in the E-base, this E-base is a base for the most conservative ranked epistemic entrenchment that respects this default theory. Intuitively, the default for *penguin*, the more specific subtype, is higher than that for the more general type *bird* because given a conflict between something's being a bird and its being a penguin, the more specific information is preferred, hence is the more highly ranked.

Now consider the question of whether birds that have wings fly. This holds if $\text{bird}(x) \wedge \text{wings}(x) \Rightarrow \text{fly}(x)$ is contained in the unique extension of the theory, or equivalently, if $\neg \text{bird}(x) \vee \neg \text{wings}(x) < \text{bird}(x) \wedge \text{wings}(x) \rightarrow \text{fly}(x)$ in the most conservative entrenchment compatible with the E-base in Figure 2. The independence assumption $\neg(\text{bird}(x) \Rightarrow \neg \text{wings}(x))$ is consistent with the theory, so by **(B7)**, it follows that $\text{bird}(x) \wedge \text{wings}(x) \Rightarrow \text{fly}(x)$. Alternatively, the rule for most conservative entrenchments says to rank formulas as low as possible consistent with the postulates (RE1)–(RE3). So the formula $\neg \text{bird}(x) \vee \neg \text{wings}(x)$ is ranked along with the nonbelief $\neg \text{bird}(x)$, whereas the formula $\text{bird}(x) \wedge \text{wings}(x) \rightarrow \text{fly}(x)$ must be ranked at least as highly as $\text{bird}(x) \rightarrow \text{fly}(x)$ because $\text{bird}(x) \rightarrow \text{fly}(x)$ implies $\text{bird}(x) \wedge \text{wings}(x) \rightarrow \text{fly}(x)$. This means that $\neg \text{bird}(x) \vee \neg \text{wings}(x) < \text{bird}(x) \wedge \text{wings}(x) \rightarrow \text{fly}(x)$, so $\text{bird}(x) \wedge \text{wings}(x) \Rightarrow \text{fly}(x)$.

However, the problem with the above definitions is that they do not give all the results that are intuitively desirable. For example, consider the above theory with the addition of the default $\text{bird}(x) \Rightarrow \text{wings}(x)$. The question now is whether penguins have wings. Since the theory contains $\text{bird}(x) \Rightarrow \neg \text{penguin}(x)$, the independence assumption that would be used to derive this conclusion, $\neg(\text{bird}(x) \Rightarrow \neg \text{penguin}(x))$, is inconsistent. In fact, there is an independence assumption consistent with the theory, and that is $\neg(\text{penguin}(x) \Rightarrow \text{wings}(x))$. This “assumption,” which is intended to be used to infer properties of nonwinged penguins from the properties of penguins, now answers the original question for us in the negative! Equivalently, the formula $\text{penguin}(x) \rightarrow \text{wings}(x)$ must be ranked at the level of $\text{bird}(x) \rightarrow \text{fly}(x)$, so that $\neg \text{penguin}(x)$ is ranked the same as $\text{penguin}(x) \rightarrow \text{wings}(x)$. This same result is a well known problem with Pearl's System Z (see [52]). The rule in force can be summed up in the statement that once an object or subclass is exceptional with respect to one property (such as flying) it is exceptional with respect to all properties (including having wings).

In our view, the problem is not with the inference patterns, but with the initial information. We consider that the conclusion that penguins have wings (which is, after all, the true state of affairs) does not follow from the initial assumptions, but rather relies on knowing the absence of a connection between not flying and having wings. To see this, note that it is possible to construct another example in which System Z gives the correct conclusion: e.g., with the default that birds eat insects, System Z infers that penguins don't eat insects. We thus consider the initial problem to be underspecified and that there is no purely formal way to derive the desired conclusions without further information. That is, it is not sufficient to use rules such as (i) that exceptions are exceptions to all properties, or (ii) that exceptions are exceptions to the smallest set of properties, but rather each property must be considered separately. A similar conclusion was reached by Goldszmidt and Pearl [31]. Thus from now on, we consider only those default theories in which such an ambiguity about the status of a property is absent. That is, we require the initial specification of a problem in default reasoning to come with a total pre-order on defaults, which corresponds to an E-base. A further complication is that the above definitions work only with propositional default theories: in practice, problems in default reasoning need to be formalized using at least a first order language.

6.2 Nonmonotonic reasoning and belief revision We have shown a connection between the unique extension of a default theory expressed in conditional logic and the most conservative belief state respecting the theory. However, from the point of view of nonmonotonic reasoning, this definition gives intuitively acceptable conclusions only with propositional default theories: i.e., if a default theory is a consistent collection of ground instances of a collection of defaults. We would like a system to operate directly with defaults so a modified revision operation must be used. The unacceptable behavior resulting from the use of most conservative belief states is easily seen with the Nixon diamond, formalized as the E-base consisting of just two equally ranked formulas $\forall x(\text{quaker}(x) \rightarrow \text{pacifist}(x))$ and $\forall x(\text{republican}(x) \rightarrow \neg \text{pacifist}(x))$. To find out whether or not Nixon, who is both a quaker and a republican, is a pacifist, we revise by $\text{quaker}(\text{nixon}) \wedge \text{republican}(\text{nixon})$. Now according to the most conservative entrenchment generated by the base, the resulting theory contains $\text{quaker}(\text{nixon}) \wedge \text{republican}(\text{nixon})$ and its consequences, as expected, but nothing else. Both rules are removed because $\neg \text{quaker}(\text{nixon}) \vee \neg \text{republican}(\text{nixon})$ is ranked the same as each rule: i.e., a default is rejected by its having just one exception!

We propose to overcome this problem by making two additional assumptions: uniqueness of names and independence of default instances. First, uniqueness of names means that any two distinct terms in the logical language denote distinct objects in the domains. Formally, this means restricting our attention to the Herbrand models of the belief language. Uniqueness of names has been suggested by Lifschitz [41] as being an intuitively desirable property of default reasoning. Given that the instances of a default all apply to different objects in the domain by uniqueness of names, the second assumption means that all these beliefs are independent. That is, there is no epistemic relationship of justification between two instances of the same default, so that when one instance of a default is removed from a theory by a revision,

all other instances of the default remain in the theory. In addition, we will use the following restricted language of defaults, which proves to be adequate for representing many problems in nonmonotonic reasoning.

Definition 6.11 A *default* is a formula of the form $\forall x \delta(x)$ where δ is a clause, each of whose literals is either of the form $[\neg]p(x)$ for some predicate symbol p and variable x or of the form $x = t$ for some ground term t . Note that the variable x must be the same in all literals in the clause.

Definition 6.12 A *default theory* is a finite, consistent ranked E-base over the language of defaults and instances of defaults.

Uniqueness of names can be formalized by adopting the standard axioms of equality and inequality into our logical language. Note, however, that an infinite Herbrand universe is assumed so that formulas such as $\forall x((x = t_1) \vee \dots \vee (x = t_n))$, usually used to express domain closure, become inconsistent. This is because it cannot be assumed that the constants used in expressing the initial defaults are all the constants in the language: to allow revision by $\neg\delta(t)$ where t is a new constant, such domain closure formulas must be inconsistent.

Independence of default instances can be formalized as follows. The idea is that on revising to accept a ground exception $\neg\epsilon(t)$ to a default ϵ , the revision procedure should leave unchanged all instances of a default δ other than $\delta(t)$. More precisely, given a ranked epistemic entrenchment Δ corresponding to a default theory, let K be the set of ground instances of the formulas in Δ and let K_t be the elements of K in which the term in the formula is t . Then to capture independence of default instances, the set of ground formulas in the revision of Δ to accept $\neg\epsilon(t)$ should be $(K_t)_{\neg\epsilon(t)}^* \cup \{K_u : u \neq t\}$. Moreover, the formulas in K_u where $u \neq t$ should be ranked as in K , whereas the ranks of the formulas in K_t should be as determined by the AGM revision operation using minimal change. The way this assumption is implemented is that on a revision by an exception $\neg\epsilon(t)$ to a default ϵ , for any default δ in the initial theory, the revised theory entails the weakened default $\forall x((x \neq t) \rightarrow \delta(x))$ and hence all the instances of δ except $\delta(t)$.

We now present an algorithm **revise**($\Gamma, \neg\epsilon(t), newrank$) for revising a belief base Γ by a formula $\neg\epsilon(t)$ with rank *newrank*, according to the most conservative entrenchment applied to the set of ground formulas $\delta(t)$ and under the assumption of independence of default instances.

revise($\Gamma, \neg\epsilon(t), newrank$)

$oldrank_{\neg\epsilon(t)} := \mathbf{rank}(\Gamma, \neg\epsilon(t))$

 if **theorem**($\epsilon(t)$) or **theorem**($\neg\epsilon(t)$) or ($newrank \leq oldrank_{\neg\epsilon(t)}$)

 return(Γ)

 else

$\Delta := \Gamma$

$oldrank_{\epsilon(t)} := \mathbf{rank}(\Gamma, \epsilon(t))$

 for each $\delta \in \Gamma$ of rank r such that $r \leq oldrank_{\epsilon(t)}$ do

 if δ covers an instance $\delta(t)$ and not(**proof**($\Delta, oldrank_{\epsilon(t)} + 1, \epsilon(t) \vee \delta(t)$))

$\Delta := \mathbf{delete}(\Delta, \delta)$

 if δ is not identical to $\delta(t)$

 let $\delta'(X) \equiv (X \neq t) \rightarrow \delta(X)$


```

50: All(X) [heavy(X) -> table(X)]           # Default
100* ~table(a)                               # Exception
? table(b)                                    # Query
>>>yes : rank = 50

# Problem A2, irrelevant information
100: red(b)                                   # Fact
? table(b)                                    # Query
>>>yes : rank = 50

# Problem A5, open domain
? All(X) [(X!=a & heavy(X)) -> table(X)]    # Query
>>>yes : rank = 50

??
>>>Complete database:
>>>100 : heavy(a)
>>>100 : heavy(b)
>>>100 : ~table(a)
>>>100 : red(b)
>>>50 : All(X) [(X!=a) -> (heavy(X) -> table(X))]

```

Note that when a revision is made to accept an exceptional instance of a default, the default is weakened to cover all but the exceptional case.

The example involving several defaults (Lifschitz's problem A3) cannot be handled by our belief revision approach to nonmonotonic reasoning because the ordering on defaults in the E-base is assumed to be a total pre-order.

```

Assumptions: Blocks A and B are heavy.
              Heavy blocks are normally located on the table.
              Heavy blocks are normally red.
              A is not on the table.
              B is not red.
Conclusions:  B is on the table.
              A is red.

```

According to our theory, any exception to a default at rank r is also an exception to all defaults at rank less than or equal to r , unless the contrary is explicitly stated, e.g., by placing an instantiation of a default at a higher rank than the default itself. So if the default that heavy blocks are on the table is placed at a higher rank than the default that blocks are red, it is concluded that B is on the table, but not that A is red. If the ranks of the rules are swapped, the opposite results are obtained. The third possibility is that the defaults are given equal rank, intuitively the most reasonable choice, yet this yields neither of the desired conclusions, as A and B become exceptions to both rules. To state that the two defaults are unrelated would require a partial order on defaults along the lines proposed by Geffner [26], but then the revision operation would not satisfy rational monotony, which corresponds to (\mathbf{K}^*8) and is thus regarded as an acceptable principle of belief revision. We regard problem A3 as un-

derspecified: the generated conclusions do not reflect a general inference pattern in nonmonotonic reasoning, but are dependent on the facts in the example, in particular the assumption that being on the table has nothing to do with being red. We conjecture that similar examples might give different intuitions.⁵

The last example in this section is Lifschitz's problem A8 on reasoning about unknown exceptions.

Assumptions: Block *A* is heavy.
 Heavy blocks are normally located on the table.
 At least one heavy block is not on the table.

Conclusion: *A* is on the table.

Note that this example is outside the scope of our theory because of the existential quantifier: the E-base representing the theory in the system contains the formula $heavy(c) \wedge \neg table(c)$ where c is a Skolem constant. The question is whether c is different from *A*. Our system does not assume that uniqueness of names applies to Skolem constants and so does not generate this conclusion. That is, we do not rule out the possibility that the block not on the table is *A*.

6.3.2 Inheritance Reasoning about inheritance networks is a standard problem in nonmonotonic reasoning, and it is relatively straightforward to solve all of Lifschitz's problems B1–B4. We present our solutions to problem set B2, tree-structured inheritance.

Assumptions: Animals normally don't fly.
 Birds are animals.
 Birds normally fly.
 Bats are animals.
 Bats normally fly.
 Emus are birds.
 Emus don't fly.

Conclusions: Animals other than birds and bats do not fly.
 Birds other than emus fly.
 Bats fly.
 Emus don't fly.

In this example, we use the rank of a formula to encode the relative strengths of the defaults. The defaults which have no exceptions are given the highest rank. Then the default rules are ordered such that the more specific defaults (that is, exceptions to other defaults) override the more general defaults.

```

100: bird(X) -> animal(X)      # Facts
100: bat(X) -> animal(X)
100: emu(X) -> bird(X)

70: emu(X) -> ~fly(X)         # Defaults
60: bird(X) -> fly(X)
60: bat(X) -> fly(X)
50: animal(X) -> ~fly(X)

```

? All(X) [animal(X) & ~bird(X) & ~bat(X) -> ~fly(X)]

>>>yes : rank = 50

? All(X) [bird(X) & ~emu(X) -> fly(X)]

>>>yes : rank = 60

? All(X) [bat(X) -> fly(X)]

>>>yes : rank = 60

? All(X) [emu(X) -> ~fly(X)]

>>>yes : rank = 70

6.3.3 Uniqueness of names In this category, Lifschitz [41] includes problems where the uniqueness of names is a default which may have exceptions, such as the following, problem C1.

Assumptions: Different names normally denote different objects.

The names ‘Ray’ and ‘Reiter’ denote the same person.

The names ‘Drew’ and ‘McDermott’ denote the same person.

Conclusion: The names ‘Ray’ and ‘Drew’ denote different people.

Despite the fact that our system employs uniqueness of names, it cannot be used to solve such problems. This is because the uniqueness of names is built into the system: there can be no exceptions since uniqueness of names by default (the first assumption above) cannot be expressed as a formula in a default theory expressed using first order logic. This assumption has the character of a meta-level inference rule, but an inference rule allowing exceptions cannot be added to the system because in the AGM approach, the base logic of beliefs is assumed to be monotonic: any formula whose negation follows from the empty theory holds only in the inconsistent belief state.

6.3.4 Reasoning about action Peppas, Foo and Wobcke [53] developed a theory of actions in which an event is modeled as a belief revision function from complete theories to complete theories, identifying an event with a revision to accept its post-condition. In a similar manner, the system can “track” the effects of multiple actions by making successive revisions, although the use of complete theories is not always necessary. The following example is Lifschitz’s problem D3.

Assumptions: After an action is performed, things normally remain as they were.

A block is on the table if and only if it is not on the floor.

When a robot grasps a block, the block will normally be in the hand.

When the robot moves a block onto the table, the block will normally be on the table.

Moving a block that is not in the hand is an exception to this rule.

Initially block A is not in the hand.

Initially block A is on the floor.

Conclusion: After the robot grasps block A, waits, and then moves it to the table, the block will not be on the floor.

For this example, there are two kinds of facts: those which may change over time, and those which are time invariant. In contrast to the examples concerning nonmonotonic reasoning in a static world, the defaults, which are time invariant, are given a higher rank than the specific facts because a change in the world is presumed to affect the facts rather than override a default. Any domain constraints have the highest rank of all.

```

100: All(X)[table(X) <-> ~ floor(X)]           # Domain Constraint

All(X)[~ holding(X) & move(X) -> ~ table(X)] # Defaults
50: All(X)[grasp(X) -> holding(X)]
50: All(X)[move(X) -> table(X)]
10: ~ holding(a)                                # Facts
10: floor(a)

10* grasp(a)                                    # Actions
10* true
10* move(a)
? ~ floor(a)                                    # Query
>>>yes : rank = 10

```

Note that the condition giving an exception to a temporal default “moving a block that is not in the hand is an exception to the previous rule” cannot be expressed in first order logic because it concerns an exception to a normal revision. This must be captured by a default stating explicitly that the outcome of the exceptional action is the negation of the expected outcome: we cannot capture chronological ignorance, only chronological denial. Note also that the waiting action is presumed to have the postcondition *true* and so is modeled by a trivial revision. Also because the revision system can be used only in a “forwards” direction, it cannot be used directly for temporal explanation or for reasoning about the unknown order of actions.

6.3.5 Autoepistemic reasoning Autoepistemic reasoning, as in Moore [46], concerns the reasoning of an agent about its own beliefs, assuming that an agent has perfect introspective ability about what it believes and what it does not believe. This is illustrated in Lifschitz’s problem E2.

```

Assumption:  At least one of the blocks A, B is on the table.
Conclusions: It is not known whether A is on the table.
             It is not known whether B is on the table.

```

These sorts of problems are simply solved if the assumptions and conclusions are converted to first order logic and the initial theory and queries taken to be beliefs of the system.

```

10: table(a) | table (b)           # Assumption
? table(a)
>>>no
? table(b)
>>>no

```

However, handling autoepistemic reasoning properly would require the base logic of beliefs to be the modal logic KD45 with an explicit belief operator.

7 Conclusion We have developed an approach to the dynamics of epistemic states which generalizes in a straightforward manner the AGM approach to belief change. Our theory, based on interpreting the notion of minimal change to apply not only to the contents of a belief state but also to its entrenchment, is characterized both axiomatically, using extensions of the AGM postulates, and constructively, using Grove's systems of spheres. A conditional logic of belief revision allowing a limited range of nested conditionals was also presented. We developed a computational approach to the dynamics of epistemic states, addressing the issue of incompleteness in the specification of an epistemic entrenchment, which we claimed to be an essential feature of problems in nonmonotonic reasoning, and gave algorithms for the computation of our theory change operations on epistemic states. Finally, we considered the connections between nonmonotonic reasoning and both conditional logic and belief revision. We showed that the unique extension of a default theory expressed in conditional logic corresponds to the most conservative epistemic state which respects the theory and presented a modified belief revision algorithm suitable for nonmonotonic reasoning. These algorithms have formed the basis of a computer system implemented using the base language of first order logic with equality in which many problems in nonmonotonic reasoning can be expressed. The system correctly handles a wide range of benchmark problems in the field.

In closing, we should emphasize that a number of assumptions have been made in this work which are summarized here for reference. First, the AGM approach to belief revision has been adopted, but in our treatment of the triviality result, following Levi [39] and Morreau [47], we have insisted on an epistemological difference between conditional and nonconditional beliefs. Second, in our construction of theory change operations, a generalized principle of minimal change applying to belief states was assumed, together with a commitment to the principle of irrelevance of syntax. Third, a partially specified ranked epistemic entrenchment was taken as standing for the most conservative entrenchment that extended it, embodying a version of foundationalism. Fourth, in nonmonotonic reasoning, both the built-in uniqueness of names and the independence of the instances of defaults were assumed. We do not mean to imply that these assumptions are the only ones possible in constructing a theory of belief revision, a semantics for indicative conditionals or a system for nonmonotonic reasoning: the overall methodology here is pluralistic in that different systems may be suitable in different contexts or in different domains. For example, one might want to allow uniqueness of names to be defeasible or one might object to **(K*8)** in belief revision or its counterpart of rational monotony in nonmonotonic reasoning. However, the most robust feature of this work is the unified approach to belief revision, the logic of indicative conditionals and nonmonotonic reasoning.

Appendix: Proofs of Theorems

Theorem 3.3 *If $K \neq K_{\perp}$, $K_{A,\alpha}^+$ is the AGM belief set K_A^+ with ranking function $rank_{A,\alpha}^+$ defined by $rank_{A,\alpha}^+(B) = \max(rank(B), \min(rank(A \rightarrow B), \alpha))$ if $B \in K_A^+$. Otherwise $K_{A,\alpha}^+ = K_{\perp}$.*

Proof: First note that our postulates (\mathbf{K}^+1) – (\mathbf{K}^+6) imply the AGM postulates (\mathbf{K}^+1) – (\mathbf{K}^+6) , so that the set of beliefs contained in $K_{A,\alpha}^+$ is the AGM set K_A^+ by (\mathbf{K}^+6) . Define the epistemic state K' using the ranking function $rank_{A,\alpha}^+$ on K_A^+ by setting $rank_{A,\alpha}^+(B) = \max(rank(B), \min(rank(A \rightarrow B), \alpha))$: this clearly satisfies (RE1) and (RE2) and so satisfies (\mathbf{K}^+1) . By definition, $rank_{A,\alpha}^+(A) \geq \alpha$ so (\mathbf{K}^+2) is satisfied. Also by definition, (\mathbf{K}^+3) and (\mathbf{K}^+4) are satisfied (that is, if $rank(A) \geq \alpha$, $rank_{A,\alpha}^+(B) = rank(B)$ for all B since $rank(B) \geq \min(rank(A \rightarrow B), rank(A))$). It is also easy to see that (\mathbf{K}^+5) is satisfied. Finally, we must show that (\mathbf{K}^+6) holds, i.e., that this definition of $rank_{A,\alpha}^+$ gives us the most conservative belief state satisfying (\mathbf{K}^+1) – (\mathbf{K}^+5) . Take any other state K' with ranking function $rank'$ satisfying these postulates. Then the belief set of this state is K_A^+ whereas for some formula $B \in K_A^+$, $rank'(B) < rank_{A,\alpha}^+(B)$. But by (\mathbf{K}^+2) , $rank'(A) \geq \alpha$ and by (\mathbf{K}^+3) , $rank'(A \rightarrow B) \geq rank(A \rightarrow B)$, so by (RE2) $rank'(B) \geq \min(rank(A \rightarrow B), \alpha) \geq rank_{A,\alpha}^+(B)$, a contradiction.

Theorem 3.4 *If $K \neq K_\perp$, K_A^- is the AGM belief set K_A^- with ranking function $rank_A^-$ defined by $rank_A^-(B) = rank(B)$ if $B \in K_A^-$. Otherwise $K_A^- = K_\perp$.*

Proof: It is straightforward to verify that (\mathbf{K}^-1) – (\mathbf{K}^-4) and (\mathbf{K}^-6) – (\mathbf{K}^-8) are satisfied by the belief state K_A^- as defined. For the recovery postulate (\mathbf{K}^-5) , consider $(K_A^-)_{A,\alpha}^+$. First $K \subseteq (K_A^-)_{A,\alpha}^+$ by the equivalent AGM postulate. Now by Theorem 3.3, the rank of a formula B in $(K_A^-)_{A,\alpha}^+ = \max(rank_A^-(B), \min(rank_A^-(A \rightarrow B), \alpha)) \geq \min(rank(A \rightarrow B), rank(A)) = rank(B)$. Thus $K \subseteq (K_A^-)_{A,\alpha}^+$. For (\mathbf{K}^-9) , suppose a more conservative state K' with ranking function $rank'$ satisfies the postulates (\mathbf{K}^-1) – (\mathbf{K}^-8) . The contents of K' must contain all formulas B such that $rank(A) < rank(A \vee B)$, hence K' must contain only these formulas if it is to satisfy (\mathbf{K}^-9) . So K' has the same content as K_A^- but differs in ranking from our K_A^- as defined. Suppose $rank'(B) < rank_A^-(B)$. In order to satisfy recovery, $rank'(A \rightarrow B)$ must equal $rank(A \rightarrow B)$, and similarly, since $B \in K_A^-$, $rank(A) < rank(A \vee B)$, and so $rank'(A \vee B)$ must also equal $rank(A \vee B)$ to satisfy recovery. But then by (RE2), $rank'(B) = rank(B)$, a contradiction.

Theorem 3.5 *If $K \neq K_\perp$, $K_{A,\alpha}^*$ is the AGM belief set K_A^* with ranking function $rank_{A,\alpha}^*$ defined by $rank_{A,\alpha}^*(B) = \max(rank(B), \min(rank(A \rightarrow B), \alpha))$ if $B \in K_A^*$. Otherwise $K_{A,\alpha}^* = K_\perp$.*

Proof: Again it is straightforward to verify (\mathbf{K}^*1) – (\mathbf{K}^*9) . Suppose $\not\vdash \neg A$ and that some epistemic state K' with ranking function $rank'$ more conservative than $K_{A,\alpha}^*$ as defined satisfies (\mathbf{K}^*1) – (\mathbf{K}^*9) . If $\neg A \notin K$ then (\mathbf{K}^*3) and (\mathbf{K}^*4) imply that both K' and $K_{A,\alpha}^*$ are identical to $K_{A,\alpha}^+$, so suppose further that $\neg A \in K$. The contents of K' must contain all formulas B such that $rank(\neg A) < rank(A \rightarrow B)$: to satisfy (\mathbf{K}^*9) , K' must contain all formulas in K of the form $A \vee B$ so that B follows when $\neg A$ is added, and K' must contain a formula of the form $A \rightarrow B$ ranked higher than $\neg A$ so that this formula's rank is preserved on adding $\neg A$. Thus K' must have the same content as $K_{A,\alpha}^*$, so suppose B is such that $rank'(B) < rank_{A,\alpha}^*(B)$. In order to satisfy (\mathbf{K}^*9) , $rank'(A \vee B)$ must equal $rank_{A,\alpha}^*(A \vee B)$ and similarly $rank'(A \rightarrow B)$ must equal $rank_{A,\alpha}^*(A \rightarrow B)$. Thus $rank'(B) = rank_{A,\alpha}^*(B)$, a contradiction.

Theorem 3.6 *If rank is a function determining a ranked system of spheres, the operation on ranked systems of spheres of complete theories defined by $rank_{A,\alpha}^*$, as follows, characterizes revision operations on consistent belief states that satisfy (\mathbf{K}^*1) – (\mathbf{K}^*10) .*

$$rank_{A,\alpha}^*(T) = \begin{cases} 0 & \text{if } A \in T \text{ and } T \in S_A \\ rank(T) & \text{if } A \in T \text{ and } T \notin S_A \\ \max(rank(T), \alpha) & \text{if } \neg A \in T. \end{cases}$$

Proof: From Grove's result, the belief set is the set of formulas contained in all theories ranked 0 and Theorem 3.5 provides the ranks of all the formulas in $K_{A,\alpha}^*$. It remains to check that for each formula B , the lowest ranked (by $rank_{A,\alpha}^*$) complete theory containing $\neg B$ is at $rank_{A,\alpha}^*(B)$. We do this first for the formulas $A \rightarrow B$ and $A \vee B$. If $A \rightarrow B \in K_{A,\alpha}^*$, the new rank of $A \rightarrow B$ is equal to the old rank of $A \rightarrow B$, and since by definition the ranking of the complete theories containing $A \wedge \neg B$ does not change, the lowest ranked complete theory containing $A \wedge \neg B$ will be ranked by $rank_{A,\alpha}^*$ the same as by $rank$. Similarly, the ranks of the theories containing $\neg A$ are moved to rank α if originally they were ranked less than α , so the lowest ranked theory containing $\neg A \wedge \neg B$ will be at $\max(rank(A \vee B), \alpha)$, as required. Hence the lowest ranked theory containing $\neg B$ is ranked at $\min(rank(B), \max(rank(A \vee B), \alpha))$, as required.

Theorem 3.7 *If rank is a function determining a ranked system of spheres, the operation on ranked systems of spheres of complete theories defined by $rank_A^-$, as follows, characterizes contraction operations on consistent belief states that satisfy (\mathbf{K}^-1) – (\mathbf{K}^-9) .*

$$rank_A^-(T) = \begin{cases} 0 & \text{if } \neg A \in T \text{ and } T \in S_{\neg A} \\ rank(T) & \text{otherwise.} \end{cases}$$

Proof: From Grove's result and the Harper identity, the belief set is given by the set of formulas contained in all theories ranked 0 and Theorem 3.4 implies that the ranks of formulas in K_A^- are unchanged from K . But a theory is ranked by $rank_A^-$ the same as by $rank$ unless it contains $\neg A$ and is contained in the smallest sphere containing a theory containing $\neg A$. Thus the ranks of formulas of the form $A \rightarrow B$ are preserved, and similarly the ranks of formulas of the form $A \vee B$ retained in K_A^- are preserved. So B has the correct rank assigned by $rank_A^-$.

Theorem 4.13 *For every faithful, consistent set Γ of BR sentences bounded to degree n , there is a belief revision model which satisfies all and only the formulas of Γ .*

Proof: Take such a set of sentences bounded to degree n and consider $\Gamma(A_1, \dots, A_{n-i})$, the set of formulas B such that $A_1 \Rightarrow (\dots A_{n-i} \Rightarrow B \dots) \in \Gamma$. We define a tree of revision sequences for each such Γ_i by induction on i , i.e., starting from the end states of the desired sequences. For $i = 0$, take all belief states τ_i whose contents agree with the base formulas of Γ_i . Clearly for such states and base formulas B , $\tau_i \models B$ iff $B \in \Gamma_i$. For $i = j + 1$, suppose by induction that for each Γ_j , there is a tree of sequences of length $\leq j$ with root τ_j such that for every conditional formula B ,

$\tau_j \models B$ iff $B \in \Gamma_j$: we need to construct a tree of sequences of length i with the same property. The root of the tree is the state K_i with contents as given by the base formulas in Γ_i and a ranking such that $A \leq B$ iff $\vdash B$ or $\neg((\neg A \vee \neg B) \Rightarrow A) \in \Gamma_i$, where A and B are base formulas. The boundedness assumption ensures that such a ranking exists (otherwise it might be possible to have two beliefs with infinitely many levels of ranking between them). To define the function $*$ at τ_i , for each base formula A , take the set $(\Gamma_i)_A^*$ of formulas B such that $A \Rightarrow B \in \Gamma_i$ and define $*(\tau_i, A)$ to be the tree τ_j corresponding to $(\Gamma_i)_A^*$: each such set is faithful by **(B14)**. Note that condition (i) on $*$ is satisfied because of **(RCEA)**, condition (ii) because of faithfulness and condition (iii) because of **(B11)** and **(B12)**. Clearly by definition, $\tau_i \models B$ iff $B \in \Gamma_i$ for a conditional formula B : for belief formulas, **(B5)** ensures the validity of the inductive step for $B_1 \wedge B_2$ and **(B13)** ensures the validity of the inductive steps for $B_1 \vee B_2$, $\neg B_1$ and $B_1 \rightarrow B_2$. Hence for all belief formulas, $\tau_i \models B$ iff $B \in \Gamma_i$. It remains to show that the transition from K_i to the root K_j of τ_j is a belief revision operation. It suffices to show that the entrenchments represented by $(\Gamma_i)_A^*$ are those that hold at $(K_i)_{A,\alpha}^*$ for some α as determined by Theorem 3.5. Now **(B9)** implies that for all base formulas $A \rightarrow B$ and $A \rightarrow C$, the relative entrenchment of $A \rightarrow B$ and $A \rightarrow C$ is preserved in K_j , and **(B10)** implies that the relative entrenchments of any base formulas $A \vee B$ and $A \vee C$ are preserved whenever these formulas are ranked higher than A . Hence the entrenchments of all the base formulas of K_j are as determined by $(\Gamma_i)_A^*$. Hence, we have constructed a belief revision model τ_n for the original set Γ such that for all belief formulas B , $\tau_n \models B$ iff $B \in \Gamma$.

Corollary 4.14 *BR is sound and complete with respect to the belief revision models.*

Proof: Soundness is easy to check. For completeness, by Lindenbaum's lemma, any nontheorem A of **BR** is contained in a bounded, faithful, consistent set of **BR** sentences containing $\neg(\text{true} \Rightarrow A)$, so by Theorem 4.13, there is a **BR** model in which A is not true.

Corollary 4.15 *There is a nontrivial belief revision model.*

Proof: We require three pairwise disjoint sentences and a belief revision model whose initial state is consistent with all three sentences. Let us take a propositional language with exactly two atoms p and q . Then the sentences $p \wedge q$, $p \wedge \neg q$ and $\neg p \wedge q$ are clearly pairwise disjoint. By completeness, there is a belief revision model whose initial state satisfies all and only the theorems of **BR**: this state is consistent with all three formulas.

Theorem 5.4 *The E-base Γ is a base for a ranked epistemic entrenchment K iff K is defined by setting $\text{rank}(A) = \max(\{r \mid \overline{\Gamma}_r \vdash A\})$ when $A \in \text{Cn}(\Gamma)$ and $\text{rank}(A) = 0$ when $A \notin \text{Cn}(\Gamma)$.*

Proof: We first verify that this definition of rank provides a ranked epistemic entrenchment. For **(RE1)**, if $A \vdash B$, any proof of A from any $\overline{\Gamma}_r$ counts as a proof of B , so the maximum over such r is a lower bound on the rank assigned to B , as required. For **(RE2)**, any two proofs of A and B can be combined to give a proof of $A \wedge B$, so the minimum of the ranks of A and B provides a lower bound on the rank of $A \wedge B$. But by **(RE1)**, the rank of $A \wedge B$ must be at most the ranks of A and B , so the minimum of the ranks of A and B provides an upper bound as well. For **(RE3)**, if $A \notin \text{Cn}(\Gamma)$ then

A is ranked 0 by definition. Now we show that the definition gives the most conservative entrenchment compatible with the E-base. Consider any proof of A from formulas in Γ . By compactness, there must be a finite set of formulas $\mathcal{A} = \{A_1, \dots, A_n\}$ such that $A \in \text{Cn}(\mathcal{A})$ but A is not a consequence of any proper subset of \mathcal{A} . Then by (RE1), $\text{rank}(A_1 \wedge \dots \wedge A_n) \leq \text{rank}(A)$, so if A_i is the lowest ranked formula in \mathcal{A} , $\text{rank}(A_i) \leq \text{rank}(A)$ by (RE2), and $\text{rank}(A_i)$ is equal to the largest r such that $A_i \in \overline{\Gamma}_r$. Thus each r such that $\overline{\Gamma}_r \vdash A$ is a lower bound on the rank of A , and so $\max(\{r \mid \overline{\Gamma}_r \vdash A\})$ is a lower bound on the rank of A . Since in the defined entrenchment the ranks of all formulas are set at this lower bound, the entrenchment defined is the most conservative entrenchment compatible with the E-base.

Corollary 5.6 *The most conservative ranked epistemic entrenchment compatible with an E-base Γ is represented by the ranked system of spheres defined by setting $\text{rank}(T)$ to be the smallest r such that T contains A for all $A \in \Gamma$ ranked greater than r .*

Proof: Take any consistent formula A of rank r . We have to show that the smallest sphere containing a complete theory containing $\neg A$ is also ranked r . By Theorem 5.4, $r = \max(\{i \mid \overline{\Gamma}_i \vdash A\})$. By definition, every complete theory ranked less than r satisfies all formulas in $\overline{\Gamma}_r$ and hence satisfies A , so the smallest sphere containing a theory containing $\neg A$ must be ranked at least r . But since the set of formulas ranked higher than r is consistent with $\neg A$ (else there would be a proof of A from this set), there must be a complete theory containing $\neg A$ ranked at r . Hence r is the rank of the smallest sphere containing a complete theory containing $\neg A$.

Theorem 5.7 *If Γ is a base for $K \neq K_\perp$, $\text{expand}(\Gamma, A, \alpha)$ is a base for $K_{A,\alpha}^+$.*

Proof: Since the algorithm adds only A and removes only formulas that can be proven in the new base with the same or higher rank, the belief set is $\text{Cn}(K \cup \{A\})$ as required. We show that the algorithm provides the correct rankings for the formulas in K_A^+ . By Theorem 3.3, a formula B is to be ranked at $\max(\text{rank}(B), \min(\text{rank}(A \rightarrow B), \alpha))$. Any proof of B in K can be converted to a proof in K_A^+ by replacing any deleted formula B' by its proof, so the rank of B in the new base will be at least that in the old. Now note that the addition of A to the base has no effect on the proofs of $A \rightarrow B$, so this formula is ranked the same in K_A^+ as in K . So, since B follows from $A \rightarrow B$ and A , the rank of B in the new base is at least $\min(\text{rank}(A \rightarrow B), \alpha)$. Hence in the new base, B is ranked at least $\max(\text{rank}(B), \min(\text{rank}(A \rightarrow B), \alpha))$. To show that this is the rank of B in the new base, note that any proof of B in K_A^+ either uses A or it doesn't: if it does, then by the deduction theorem we have a proof of $A \rightarrow B$ in K , and if it doesn't, we have a proof of B in K . Hence the rank of B in the new base is $\max(\text{rank}(B), \min(\text{rank}(A \rightarrow B), \alpha))$, as required.

Theorem 5.8 *If Γ is a base for $K \neq K_\perp$, $\text{contract}(\Gamma, A)$ is a base for K_A^- .*

Proof: Consider any formula B such that $\text{rank}(A) < \text{rank}(A \vee B)$. Then the base computed by the algorithm contains $A \vee B$ since this is ranked higher than the rank of A , and also contains $A \rightarrow B$ by the test on deletion of formulas from the base, hence B follows from the new base. Conversely, if B follows from the new base then there is a proof of B from formulas B_1, \dots, B_n in the new base and hence proofs of $A \vee B_1, \dots, A \vee B_n$ in the old base, all ranked higher than the rank of A , hence a proof

of $A \vee B$ ranked higher than the rank of A . So the belief set computed by the algorithm is K_A^- as determined by the (C^-) condition. We show that the algorithm provides the correct rankings for formulas of the forms $A \vee B$ and $A \rightarrow B$. By Theorem 3.4, all formulas are to be ranked as in K . First, since $A \vee B$ is retained only if it is ranked higher than A and the algorithm does not affect the ranks of such formulas, the ranks of $A \vee B$ will be the same in K and K_A^- . Second, note that K_A^- contains $A \rightarrow B$ iff K contains $A \rightarrow B$. Then any proof of $A \rightarrow B$ in K will remain a proof in K_A^- because changing base formulas C to $A \rightarrow C$ has no effect on the ranks of proofs of these formulas, and conversely, any proof of $A \rightarrow B$ in the new base will correspond to a proof in the old base, thus preserving the rank of $A \rightarrow B$. By (RE2), the ranks of all base formulas B are as required.

Theorem 5.9 *If Γ is a base for $K \neq K_\perp$, $\text{revise}(\Gamma, A, \alpha)$ is a base for $K_{A,\alpha}^*$.*

Proof: By analogy to contraction, the belief set computed by the algorithm is K_A^* as determined by the (C^*) condition. We now show that the ranks of formulas in the new base are as required by Theorem 3.5. First, a formula of the form $A \rightarrow B$ must have the same rank in $K_{A,\alpha}^*$ as in K . This follows from the fact that the revised state contains this formula only if it is ranked higher than the old rank of A , and these formulas are unchanged by the algorithm because a proof of $A \rightarrow B$ cannot involve the added formula A . Consider formulas of the form $A \vee B$. Since changing base formulas C to $A \vee C$ does not affect the ranks of proofs of such formulas, these formulas will be ranked as in K except when they follow from A which is added at rank α . Thus the new rank of $A \vee B$ is $\max(\text{rank}(A \vee B), \alpha)$. By (RE2), the ranks of all base formulas B are correct.

Theorem 6.8 *Extensions of a BR theory Γ are in correspondence with the most conservative ranked epistemic entrenchments respecting Γ .*

Proof: *extension \Rightarrow most conservative entrenchment*

First suppose we have an extension Γ^* of a given BR theory Γ . Define a ranked entrenchment compatible with Γ^* according to the entrenchments that Γ^* represents, i.e., $A \leq B$ iff $\Box B \vee \neg((\neg A \vee \neg B) \Rightarrow A) \in \Gamma$. The finiteness of Γ ensures that this is possible. For a contradiction, suppose this entrenchment is not the most conservative entrenchment compatible with Γ . Then the rank of some B can be lowered to that of some A resulting in a more conservative entrenchment that is also compatible with Γ . This new entrenchment agrees with the original one on all formulas less entrenched than A , and also corresponds to a BR model of Γ . Let Γ' be the set of sentences satisfied in this model. Clearly Γ' contains Γ . But Γ' also contains the formula $\neg((\neg A \vee \neg B) \Rightarrow B)$ corresponding to the entrenchment $B \leq A$, which is not contained in Γ^* . It follows that Γ' is an extension of Γ , which contradicts Γ^* being an extension of Γ . To verify this, note that any independence assumption $\neg(A' \Rightarrow \neg B')$, representing the entrenchment $\neg A' = \neg A' \vee \neg B'$, which is stronger than $\neg((\neg A \vee \neg B) \Rightarrow B)$, i.e., such that $\neg A \vee \neg B \vdash A'$, is contained in Γ' iff it is contained in Γ^* , because if $\neg A' < A$, the new entrenchment agrees with the old over such formulas, and if $\neg A' = A$, the set of formulas entrenched equally as A contains that for the old entrenchment. Thus Γ^* is not an extension of Γ , a contradiction.

most conservative entrenchment \Rightarrow extension

Suppose we have the most conservative ranked epistemic entrenchment compatible

with a BR theory Γ . Take the BR model corresponding to this entrenchment, and let Γ^* be the set of formulas holding at the state satisfying Γ in this model. For a contradiction, suppose that Γ^* is not an extension of Γ . Then there is a set Γ' which contains Γ and a stronger set of independence assumptions. In particular, Γ' contains independence assumptions $\neg(A \Rightarrow \neg B)$ not contained in Γ^* . Choose one so that the entrenchment of $\neg A$ is minimized, so that for any formulas C and D less entrenched than $\neg A$, Γ' represents $C < D$ iff Γ^* represents $C < D$. Now from this independence assumption, it follows using $A \Rightarrow A$ that $\neg(A \Rightarrow (\neg A \vee \neg B))$, i.e., $\neg(((A \wedge B) \vee A) \Rightarrow (\neg A \vee \neg B))$, which corresponds to the entrenchment $\neg A \vee \neg B \leq \neg A$. Since $\neg A \leq \neg A \vee \neg B$ by (EE2), we now have $\neg A = \neg A \vee \neg B$. This means the entrenchment was not the most conservative entrenchment compatible with Γ , which is the desired contradiction.

Corollary 6.9 *A consistent default theory has a unique extension.*

Proof: Since the set of all ranked epistemic entrenchments compatible with Γ is closed under meet, there is a unique most conservative entrenchment compatible with Γ .

Corollary 6.13 *If Γ is a base for Δ , $\text{revise}(\Gamma, \neg\epsilon(t), \alpha)$ is a base for $\Delta_{-\epsilon(t), \alpha}^*$ under independence of default instances.*

Proof: By Theorem 5.9, the algorithm correctly determines the ranks of the instances of defaults $\delta(t)$. If $u \neq t$ and $\delta(u)$ is in Δ , then $\delta(u)$ is also in $\Delta_{-\epsilon(t), \alpha}^*$ with the same rank, since any proof of such a formula in the original base remains a proof in the revised base, and vice versa.

Acknowledgments This work was supported by the Australian Research Council. I would like to thank Norman Foo for introducing me to research on belief revision and for coordinating the Knowledge Systems Group at the University of Sydney, the members of the group, both past and present, for discussions over the last few years, Abhaya Nayak for discussion of some of the more philosophical aspects of this paper, and especially Simon Dixon for collaboration on the development of the theory change algorithms and on the implementation of the belief revision system.

NOTES

1. It is possible to define algorithms conforming to either intuition, but implementing Spohn's intuition is computationally more expensive.
2. For readers interested in such matters, the system is written in C and uses a variant of OL-resolution with paramodulation, (cf. Chang and Lee [5]).
3. John Slaney (personal communication, April 3, 1994) has suggested the use of a model generator to find upper bounds on the rank of A , but this procedure has not yet been implemented.
4. For an implementation of the base contraction operation of Williams [61], this step can simply be omitted; cf. Dixon [11].
5. Interestingly, in a series of psychological experiments, Elio and Pelletier [14] have found that Lifschitz's intuitions here are not supported nearly as strongly as in the other examples. One possible reason for this which is in line with our approach is that sub-

jects are trying to rank the given information, and this ranking is different across subjects/examples. Needless to say, the evidence is inconclusive.

REFERENCES

- [1] Alchourrón, C. E., P. Gärdenfors and D. Makinson, "On the logic of theory change: partial meet contraction and revision functions," *The Journal of Symbolic Logic*, vol. 50 (1985), pp. 510–530.
- [2] Arlo-Costa, H. L., and S. J. Shapiro, "Maps between nonmonotonic and conditional logic," pp. 553–564 in *Proceedings of the Third International Conference on Principles of Knowledge, Representation, and Reasoning*, Morgan Kaufmann, San Mateo, 1992.
- [3] Bell, J., "Pragmatic logics," pp. 50–60 in *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, 1991.
- [4] Boutilier, C., "Revision sequences and nested conditionals," pp. 519–525 in *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo, 1993.
- [5] Chang, C. L., and R. C. T. Lee, *Symbolic Logic and Mechanical Theorem Proving*, Academic Press, New York, 1973.
- [6] Crocco, G., and P. Lamarre, "On the connection between non-monotonic inference systems and conditional logics," pp. 565–571 in *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, 1992.
- [7] Dalal, M., "Investigations into a theory of knowledge base revision: preliminary report," pp. 475–479 in *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, Morgan Kaufmann, San Mateo, 1988.
- [8] Darwiche, A., and J. Pearl, "On the logic of iterated belief revision," pp. 5–23 in *Proceedings of the Fourth Conference on Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann, San Mateo, 1994.
- [9] Delgrande, J. P., "A first-order conditional logic for prototypical properties," *Artificial Intelligence*, vol. 33 (1987), pp. 105–130.
- [10] Delgrande, J. P., "An approach to default reasoning based on a first-order conditional logic: revised report," *Artificial Intelligence*, vol. 36 (1988), pp. 63–90.
- [11] Dixon, S. E., "A finite base belief revision system," pp. 445–451 in *Proceedings of the Sixteenth Australian Computer Science Conference*, Queensland University of Technology, Brisbane, 1993.
- [12] Dixon, S. E., and W. R. Wobcke, "The implementation of a first-order logic AGM belief revision system," pp. 40–47 in *Proceedings of the Fifth IEEE International Conference on Tools with Artificial Intelligence*, IEEE Computer Society Press, Los Alamitos, 1993.
- [13] Dixon, S. E. and W. R. Wobcke, "A nonmonotonic reasoning belief revision system," pp. 151–160 in *Proceedings of the Eighteenth Australasian Computer Science Conference*, Flinders University of South Australia, University of Adelaide, and University of South Australia, Adelaide, 1995.

- [14] Elio, R. E., and F. J. Pelletier, "Human benchmarks on AI's benchmark problems," pp. 406–411 in *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum, Hillsdale, 1993.
- [15] Etherington, D. W., *Reasoning with Incomplete Information*, Pitman, London, 1987.
- [16] Fariñas del Cerro, L., A. Herzig and J. Lang, "From ordering-based nonmonotonic reasoning to conditional logics," *Artificial Intelligence*, vol. 66 (1994), pp. 375–393.
- [17] Fuhrmann, A., "Theory contraction through base contraction," *Journal of Philosophical Logic*, vol. 20 (1991), pp. 175–203.
- [18] Gabbay, D. M., "Theoretical foundations for non-monotonic reasoning in expert systems," pp. 439–457 in *Proceedings of the NATO Advanced Study Institute on Logics and Models of Concurrent Systems*, edited by K. R. Apt, Springer-Verlag, Berlin, 1985.
- [19] Gärdenfors, P., "Belief revisions and the Ramsey Test for conditionals," *Philosophical Review*, vol. 95 (1986), pp. 81–93.
- [20] Gärdenfors, P., "Variations on the Ramsey Test: more triviality results," *Studia Logica*, vol. 46 (1987), pp. 319–325.
- [21] Gärdenfors, P., *Knowledge in Flux*, MIT Press, Cambridge, 1988.
- [22] Gärdenfors, P., "Belief revision and nonmonotonic logic: two sides of the same coin?" pp. 768–773 in *Proceedings of the Ninth European Conference on Artificial Intelligence*, Pitman, London, 1990.
- [23] Gärdenfors, P., "Nonmonotonic inferences based on expectations: a preliminary report," pp. 585–590 in *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, 1991.
- [24] Gärdenfors, P., S. Lindström, M. Morreau and W. Rabinowicz, "The negative Ramsey Test: another triviality result," pp. 129–134 in *The Logic of Theory Change*, edited by A. Fuhrmann and M. Morreau, Springer-Verlag, Berlin, 1991.
- [25] Gärdenfors, P., and D. Makinson, "Revisions of knowledge systems using epistemic entrenchment," pp. 83–95 in *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann, Los Altos, 1988.
- [26] Geffner, H. A., *Default Reasoning: Causal and Conditional Theories*, MIT Press, Cambridge, 1992.
- [27] Geffner, H. A., and J. Pearl, "Conditional entailment: bridging two approaches to default reasoning," *Artificial Intelligence*, vol. 53 (1992), pp. 209–244.
- [28] Ginsberg, M. L., and D. E. Smith, "Reasoning about action I: a possible worlds approach," *Artificial Intelligence*, vol. 35 (1988), pp. 165–195.
- [29] Goldszmidt, M., P. Morris and J. Pearl, "A maximum entropy approach to nonmonotonic reasoning," pp. 646–652 in *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, AAAI Press, Menlo Park, 1990.
- [30] Goldszmidt, M., and J. Pearl, "On the relation between rational closure and system Z," pp. 130–140 in *Proceedings of the Third International Workshop on Nonmonotonic Reasoning*.
- [31] Goldszmidt, M., and J. Pearl, "System Z⁺: a formalism for reasoning with variable-strength defaults," pp. 399–404 in *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, AAAI Press, Menlo Park, 1991.

- [32] Grove, A., "Two modellings for theory change," *Journal of Philosophical Logic*, vol. 17 (1988), pp. 157–170.
- [33] Hansson, S. O., "New operators for theory change," *Theoria*, vol. 55 (1989), pp. 114–136.
- [34] Hansson, S. O., "Belief contraction without recovery," *Studia Logica*, vol. 50 (1991), pp. 251–260.
- [35] Hansson, S. O., "A dyadic representation of belief," pp. 89–121 in *Belief Revision*, edited by P. Gärdenfors, Cambridge University Press, Cambridge, 1992.
- [36] Katsuno, H. and K. Satoh, "A unified view of consequence relation, belief revision and conditional logic," pp. 406–412 in *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo, 1991.
- [37] Kraus, S., D. Lehmann, and M. Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics," *Artificial Intelligence*, vol. 44 (1990), pp. 167–207.
- [38] Lehmann, D. and M. Magidor, "What does a conditional knowledge base entail?" *Artificial Intelligence*, vol. 55 (1992), pp. 1–60.
- [39] Levi, I., "Iteration of conditionals and the Ramsey Test," *Synthese*, vol. 76 (1988), pp. 49–81.
- [40] Lewis, D. K., *Counterfactuals*, Blackwell, Oxford, 1973.
- [41] Lifschitz, V., "Benchmark problems for formal nonmonotonic reasoning," pp. 202–219 in *Non-Monotonic Reasoning*, edited by M. Reinfrank, J. de Kleer, M. L. Ginsberg and E. Sandewall, Springer-Verlag, Berlin, 1989.
- [42] Makinson, D., "On the status of the postulate of recovery in the logic of theory change," *Journal of Philosophical Logic*, vol. 16 (1987), pp. 383–394.
- [43] Makinson, D., "General theory of cumulative inference," pp. 1–18 in *Non-Monotonic Reasoning*, edited by M. Reinfrank, J. de Kleer, M. L. Ginsberg and E. Sandewall, Springer-Verlag, Berlin, 1989.
- [44] Makinson, D., and P. Gärdenfors, "Relations between the logic of theory change and nonmonotonic logic," pp. 185–205 in *The Logic of Theory Change*, edited by A. Fuhrmann and M. Morreau, Springer-Verlag, Berlin, 1991.
- [45] McCarthy, J. M., and P. J. Hayes, "Some philosophical problems from the standpoint of artificial intelligence," pp. 463–502 in *Machine Intelligence 4*, edited by B. Meltzer and D. Michie, Edinburgh University Press, Edinburgh, 1969.
- [46] Moore, R. C., "Semantical considerations on nonmonotonic logic," *Artificial Intelligence*, vol. 25 (1985), pp. 75–94.
- [47] Morreau, M., "Epistemic semantics for counterfactuals," *Journal of Philosophical Logic*, vol. 21 (1992), pp. 33–62.
- [48] Nayak, A. C., "Foundational belief change" *Journal of Philosophical Logic*, vol. 23 (1994), pp. 495–533.
- [49] Nayak, A. C., "Iterated belief change based on epistemic entrenchment," *Erkenntnis*, vol. 41 (1994), pp. 353–390.
- [50] Nayak, A. C., N. Y. Foo, M. Pagnucco and A. Sattar, "Entrenchment kinematics 101," pp. 157–164 in *Proceedings of the Seventh Australian Joint Conference on Artificial Intelligence*, World Scientific, Singapore, 1994.

- [51] Nebel, B., "Belief revision and default reasoning: syntax-based approaches," pp. 417–428 in *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, 1991.
- [52] Pearl, J., "System Z: a natural ordering of defaults with tractable applications to non-monotonic reasoning," pp. 121–135 in *Proceedings of the Third Conference on Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann, San Mateo, 1990.
- [53] Peppas, P., N. Y. Foo and W. R. Wobcke, "Events as theory operators," pp. 413–426 in *Proceedings of the First World Conference on the Fundamentals of Artificial Intelligence*, Angkor, Paris, 1991.
- [54] Poole, D., "A logical framework for default reasoning," *Artificial Intelligence*, vol. 36 (1988), pp. 27–47.
- [55] Ramsey, F. P., "General propositions and causality," pp. 237–255 in F. P. Ramsey *The Foundations of Mathematics and Other Logical Essays*, Routledge and Kegan Paul, New York, 1931.
- [56] Rott, H., "A nonmonotonic conditional logic for belief revision," pp. 135–181 in *The Logic of Theory Change*, edited by A. Fuhrmann and M. Morreau, Springer-Verlag, Berlin, 1991.
- [57] Rott, H., "Preferential belief change using generalized epistemic entrenchment," *Journal of Logic, Language and Information*, vol. 1 (1992), pp. 45–78.
- [58] Spohn, W., "Ordinal conditional functions: a dynamic theory of epistemic states," pp. 105–134 in *Causation in Decision, Belief Change, and Statistics, II*, edited by W. L. Harper and B. Skyrms, Kluwer Academic Publishers, Dordrecht, 1988.
- [59] Stalnaker, R. C., "A theory of conditionals," pp. 98–112 in *Studies in Logical Theory*, edited by N. Rescher, Blackwell, Oxford, 1968.
- [60] van Fraassen, B. C., "Singular terms, truth-value gaps, and free logic," *Journal of Philosophy*, vol. 63 (1966), pp. 481–495.
- [61] Williams, M. A., "Two operators for theory base change," pp. 259–265 in *Proceedings of the Fifth Australian Joint Conference on Artificial Intelligence*, World Scientific, Singapore, 1992.
- [62] Williams, M. A., "Transmutations of knowledge systems," pp. 619–629 in *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, 1994.
- [63] Wobcke, W. R., "On the use of epistemic entrenchment in nonmonotonic reasoning," pp. 324–328 in *Proceedings of the Tenth European Conference on Artificial Intelligence*, Wiley, Chichester, 1992.
- [64] Wobcke, W. R., "Algorithms for iterative belief revision," pp. 165–172 in *Proceedings of the Seventh Australian Joint Conference on Artificial Intelligence*, World Scientific, Singapore, 1994.

*Knowledge Systems Group
Basser Department of Computer Science
University of Sydney, Sydney NSW 2006
Australia
email: Wobcke@cs.su.oz.au*