Reset reproduction of article published in *Belief Revision* (P. Gärdenfors, ed.), Cambridge: Cambridge University Press, 1992, pp. 29–51. Reprinted July 1994.

Reason Maintenance and Belief Revision
Foundations vs. Coherence Theories

Jon Doyle

Laboratory for Computer Science
Massachusetts Institute of Technology

In memory of Norma Charlotte Schleif Miller, 1898-1991

1 INTRODUCTION

Recent years have seen considerable work on two approaches to belief revision: the so-called *foundations* and *coherence* approaches. The foundations approach supposes that a rational agent derives its beliefs from justifications or reasons for these beliefs: in particular, that the agent holds some belief if and only if it possesses a satisfactory reason for that belief. According to the foundations approach, beliefs change as the agent adopts or abandons reasons. The coherence approach, in contrast, maintains that pedigrees do not matter for rational beliefs, but that the agent instead holds some belief just as long as it logically coheres with the agent's other beliefs. More specifically, the coherence approach supposes that revisions conform to minimal change principles and conserve as many beliefs as possible as specific beliefs are added or removed. The artificial intelligence notion of *reason maintenance system* (Doyle, 1979) (also called "truth maintenance system") has been viewed as exemplifying the foundations approach, as it explicitly computes sets of beliefs from sets of recorded reasons. The so-called AGM theory of Alchourrón, Gärdenfors and Makinson (1985; 1988) exemplifies the coherence approach with its formal postulates characterizing conservative belief revision.

Although philosophical work on the coherence approach influenced at least some of the work on the foundations approach (e.g., (Doyle, 1979) draws inspiration from (Quine, 1953; Quine and Ullian, 1978)), Harman (1986) and Gärdenfors (1990) view the two approaches as antithetical. Gärdenfors has presented perhaps the most direct argument for preferring the coherence approach to the foundations approach. He argues that the foundations approach involves excessive computational expense, that it conflicts with observed psychological behavior, and that the coherence approach subsumes the foundations approach in the sense that one can sometimes reconstruct the information contained in reasons from the information about "epistemic entrenchment" guiding conservative revision.

In this paper, we examine Gärdenfors's criticisms of the foundations approach. We argue that the coherence and foundations approaches differ less than has been supposed, in that the fundamental concerns of the coherence approach for conservatism in belief revision apply in exactly the same way in the foundations approach. We also argue that the foundations approach represents the most direct way of mechanizing the coherence approach. Moreover, the computational costs of revisions based on epistemic entrenchment appear to equal or exceed those of revisions based on reasons, in the sense that any entrenchment ordering from which information about reasons may be recovered will be at least as costly to update as the reasons it represents. We conclude that while the coherence approach offers a valuable perspective on belief revision, it does not yet provide an adequate theoretical or practical basis for characterizing or mechanizing belief revision.

2 THE COHERENCE APPROACH TO BELIEF REVISION

The coherence approach to belief revision maintains that an agent holds some belief just as long as it coheres with the agent's other beliefs, independent of how they may have been inferred or adopted. In other words, the coherence approach focuses on logical and psychological relations among beliefs rather than on inferential pedigrees. While one belief may be related to more beliefs than another, no belief is more fundamental than another. Indeed, when the set of beliefs contains sufficiently many of the consequences of these beliefs, one can usually derive any single belief from the others. A deductively closed set of beliefs represents the extreme case in which each belief follows from all the others.

One arrives at different coherence theories by choosing different ways of making the notion of "coherence" precise. Typical theories require that belief states should be *logically consistent*, and that changes of state should be *epistemologically conservative* in the sense that (roughly speaking) the agent retains as many of its beliefs as possible when it accommodates its beliefs to new information. (Quine (1970) calls epistemological conservatism "minimum mutilation"; Harman (1986) also calls it "conservativity.") Some authors, for example Harman (1986), supplement or supplant consistency with other relations of implication and explanation among beliefs, but these just indicate additional connections among beliefs rather than reasons stating why some beliefs are held. The requirement of consistency reflects a concern with the logical content of the agent's beliefs: inconsistent beliefs describe no world, and so cannot be useful. The requirement of conservatism reflects a concern with the economics of reasoning: information is valuable (costly to acquire), and so loss of information should be minimized.

A precise coherence approach must spell out just what these two requirements mean.

Logical consistency has an accepted definition, so the most pressing task is to provide a precise notion of conservatism. Two of the simplest ways of comparing the sizes of changes in beliefs compare the sets of added and subtracted beliefs or the cardinality of these sets. But measures of the size of changes need not be as simple as counting the number of beliefs adopted or abandoned.

2.1 The AGM formalization

In order to formalize the essence of conservative revision independent of particular choices of measures, Alchourrón, Gärdenfors and Makinson (1985) developed an axiomatic approach to belief revision that avoids commitment to any particular measure. (Gärdenfors (1988) treats this approach along with related materials by the authors.) We summarize their approach using an adaptation of the notations of (Alchourrón *et al.*, 1985). We suppose that $\mathcal{L}$ is a propositional language over the standard sentential connectives ($\neg$, $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$), denote individual propositions by α, β, and γ, and denote sets of propositions by K and K'. We write $\vdash$ to mean classical propositional derivability, and write Cn to mean the corresponding closure operator

$$Cn(K) \stackrel{\text{def}}{=} \{\alpha \in \mathcal{L} \mid K \vdash \alpha\}.$$

The AGM approach models states of belief by sets of propositions. In some treatments, such as that of Gärdenfors (1988), states of belief are modeled by deductively closed (but not necessarily consistent) sets of propositions, that is, propositional *theories* $K \subseteq \mathcal{L}$ such that $K = Cn(K)$. The intent is to capture the import of the agent's beliefs, not necessarily what the agent will explicitly assent to or represent. Many of the theoretical results about belief revision concern this case of closed belief states. In other treatments, however, states of belief are modeled by sets of propositions that need not be deductively closed. These sets, called *belief bases*, represent the beliefs contained in their deductive closure. Formally, we say that K' is a base for K whenever $K = Cn(K')$ (even when $K = K'$). Naturally, a given theory can be represented by many different belief bases. The case of greatest practical interest is when the belief base K' is finite (and small), but not every theory has a finite basis.

The AGM approach considers three types of operations on belief states. For each belief state K and proposition α we have:

Expansion: Expanding K with α, written $K+\alpha$, means adding α to K and requiring that the result be a (possibly inconsistent) belief state.

Contraction: Contracting K with respect to α, written $K \dot{-} \alpha$, means removing α from K in such a way to result in a belief state.

Revision: Revising K with α, written $K \dot{+} \alpha$, means adding α to K in such a way that the result is a consistent belief state.

Expansion is naturally defined in terms of the union of the set of beliefs and the new proposition. In the belief base model, we may take the expansion of K by α as this union itself

$$K + \alpha \stackrel{\text{def}}{=} K \cup \{\alpha\},$$

while for closed belief states, we take the expansion to be the closure of this union

$$K + \alpha \stackrel{\text{def}}{=} Cn(K \cup \{\alpha\}).$$

Contraction and revision, on the other hand, have no single natural definitions, only the standard requirement that the change made be as small as possible so as to minimize unnecessary loss of knowledge. This requirement does not define these operations since there are usually several ways to get rid of some belief. In the case of contraction, for example, there is generally no largest belief state $K' \subseteq K$ such that $K' \not\vdash \alpha$. For example, if α and β are logically independent, $K = \{\alpha, \beta\}$, and we wish to determine $K \dot{-} (\alpha \wedge \beta)$, neither $\{\alpha\}$ nor $\{\beta\}$ entails $\alpha \wedge \beta$, but neither is one a subset of the other.

Prevented from identifying unique natural contraction and revision operations, Alchourrón, Gärdenfors, and Makinson formulate and motivate sets of rationality postulates that these operations should satisfy. We do not need to review these postulates here, other than to mention that the postulates for contraction and revision are logically equivalent if the revision $K \dot{+} \alpha$ of a theory K is defined by means of the *Levi identity*

$$K \dot{+} \alpha \stackrel{\text{def}}{=} (K \dot{-} \neg\alpha) + \alpha,$$

so that revision by α is equivalent to contracting by $\neg\alpha$ to remove any inconsistent beliefs and then expanding with α. Alternatively, one can define contractions in terms of revisions by means of the *Harper identity*

$$K \dot{-} \alpha \stackrel{\text{def}}{=} (K \dot{+} \neg\alpha) \cap K,$$

so that the contraction by α is equivalent to taking those beliefs that would be preserved if $\neg\alpha$ were now believed.

2.2 Epistemic entrenchment

Gärdenfors (1988) views the behaviors described by the AGM postulates as arising from a more fundamental notion, that of *epistemic entrenchment*. Epistemic entrenchment is characterized by a complete preorder (a reflexive and transitive relation) over propositions which indicates which propositions are more valuable than others. This ordering influences revisions by the requirement that revisions retain more entrenched beliefs in preference to less entrenched ones. It may change over time or with the state of belief.

If α and β are propositions, we write $\alpha \leq \beta$ to mean that β is at least as epistemically entrenched as α. We write $\alpha < \beta$ (the strict part of this order) to mean that β is more entrenched than α, that is, that $\alpha \leq \beta$ and $\beta \not\leq \alpha$. We write $\alpha \sim \beta$ (the reflexive part of the order) to mean that $\alpha \leq \beta$ and $\beta \leq \alpha$. The following postulates characterize the qualitative structure of epistemic entrenchment.

($\leq$1) If $\alpha \leq \beta$ and $\beta \leq \gamma$, then $\alpha \leq \gamma$; *(transitivity)*

($\leq$2) If $\alpha \vdash \beta$, then $\alpha \leq \beta$; *(dominance)*

($\leq$3) Either $\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$; *(conjunctiveness)*

($\leq$4) If K is a consistent theory, then $\alpha \leq \beta$ for all β iff $\alpha \notin K$; *(minimality)*

($\leq$5) If $\alpha \leq \beta$ for all α, then $\vdash \beta$. *(maximality)*

Postulate ($\leq$1) just says that $\leq$ is an ordering relation, while the other postulates all concern how the logic of propositions interacts with the ordering. Postulate ($\leq$2) says that α entails β, then retracting α is a smaller change than retracting β, since the closure requirement on belief states means that β cannot be retracted without giving up α as well. Postulate ($\leq$3) reflects the fact that a conjunction cannot be retracted without giving up at least one of its conjuncts. Taken together postulates ($\leq$1)-($\leq$3) imply that $\leq$ is a complete ordering, that is, that either $x \leq \beta$ or $\beta \leq \alpha$. Propositions not in a belief state are minimally entrenched in that state, according to ($\leq$4), and according to ($\leq$5), the only way a proposition can be maximally entrenched is if it is logically valid.

The influence of epistemic entrenchment on belief revisions is characterized by two conditions relating entrenchment orderings and contraction functions over theories. The first condition,

$$\alpha \leq \beta \text{ iff either } \alpha \notin K \dot{-} (\alpha \wedge \beta) \text{ or } \vdash \alpha \wedge \beta, \tag{1}$$

says that in contracting a theory K with respect to a conjunction, we must give up the conjunct of lesser epistemic entrenchment, or both conjuncts if they are equally entrenched. It says, roughly speaking, that $\alpha < \beta$ is the same as $\beta \in K \dot{-} (\alpha \wedge \beta)$. The second condition,

$$\beta \in K \dot{-} \alpha \text{ iff } \beta \in K \text{ and either } \alpha < \alpha \vee \beta \text{ or } \vdash \alpha,$$

explicitly characterizes contraction functions in terms of epistemic entrenchment orderings. Using the contraction condition (1), Gärdenfors and Makinson (1988) prove that the postulates ($\leq$1)-($\leq$5) characterize the same class of belief revision operations as do the AGM postulates for contraction and revision.

3 THE FOUNDATIONS APPROACH

Where the coherence approach seeks to view all beliefs as held independently, related only by coherence requirements, the foundations approach divides beliefs into two classes: those justified by other beliefs, and those not justified by other beliefs. The former constitute derived beliefs, while one may view the latter as "self-justifying" beliefs or basic postulates. In contrast to the way beliefs provide each other with mutual support in the coherence approach, and so provide little help when it comes to giving explanations of why one believes something, the foundations approach provides explanations of beliefs by requiring that each derived belief be supportable by means of some noncircular arguments from basic beliefs. This noncircularity also allows one to use the arguments to determine changes to the overall set of beliefs; one should retain a belief, even after removing one of its justifications, as long as independent justifications remain, and one should abandon a belief after removing or invalidating the last of its justifications. In other words, one should hold a derived belief if and only if it has at least one noncircular argument from foundational beliefs.

3.1 Reason maintenance

Harman (1986) and Gärdenfors (1990) view reason maintenance systems as prime exemplars of the foundations approach. A reason maintenance system acts as a subsystem of a more general system for reasoning. It helps revise the database states of the overall system by using records of inferences or computations to trace the consequences of initial changes. By keeping track of what information has been computed from what, such it reconstructs the information "derivable" from given information. Although we find it convenient to think of these bits of information and derivations as beliefs and arguments, one may apply reason maintenance more generally to all sorts of computational or mental structures.

For concreteness, we follow Harman and Gärdenfors and focus on RMS, the particular reason maintenance system developed by the author (Doyle, 1979). We may formalize its essential features as follows, simplifying its complex actual structure in ways that do not matter for the present discussion. (See (Doyle, 1983a; Doyle, 1983b) for more detailed discussions.)

States of RMS contain two types of elements: *nodes* and *reasons*. We write $\mathcal{N}$ to denote the set of possible nodes, and $\mathcal{R}$ to denote the set of possible reasons. RMS uses nodes to represent information (beliefs, desires, rules, procedures, database elements, etc.) of significance to the overall reasoning system, but those "external" meanings have no bearing on the "internal" operation of RMS. RMS uses reasons to represent inferences, or more precisely, specific inference rules. Because nodes need not represent beliefs, RMS imposes no logic on nodes. Instead, the only relations among nodes are those indicated explicitly by reasons. These may, if desired, encode logical relations directly. For simplicity, we assume no nodes are reasons, so that each state consists of a set $N \subseteq \mathcal{N}$ of nodes and a set $R \subseteq \mathcal{R}$ of reasons.[1] We say that each node in N is *in* (the set of beliefs), and that each node in $\mathcal{N} \setminus N$ is *out* (of the set of beliefs).

The original RMS provided two types of reasons: *support-list* reasons and *conditional-proof* reasons. For simplicity, we will ignore conditional-proof reasons and assume that all reasons are support-list reasons. Each support-list reason takes the form (I, O, c) where $I, O \subseteq \mathcal{N}$ and $c \in \mathcal{N}$. The components I, O and c are called the *inlist*, the *outlist*, and the *consequent*, respectively. We call the reason *monotonic* if O is empty, and *nonmonotonic* otherwise. Each reason is interpreted as a rule stipulating inferences the RMS must make, according to which the consequent holds if all of the nodes in the inlist are held and none of the nodes in the outlist are held.

A state (N, R) is a *legal* state of RMS just in case N consists exactly of the *grounded* consequences of the reasons R. Formally, (N, R) is a legal state just in case

1. If $(I, O, c) \in R$, $I \subseteq N$, and $N \cap O = \emptyset$, then $c \in N$; and
2. If $n \in N$, then there is a finite sequence $\langle n_0, \ldots, n_m \rangle$ of elements of $N \cup R$ such that $n = n_m$ and for each $i \leq m$, either $n_i \in R$, or there is some $j < i$ such that
   (a) $n_j = (I, O, n_i)$,
   (b) for each $x \in I$, $x = n_k$ for some $k < j$, and
   (c) $x \notin N$ for each $x \in O$.

In other words, (N, R) is a legal state if the set of nodes *satisfies* every reason in R and if every node in N is supported by a noncircular argument from the *valid* reasons in R. Each set of reasons supports 0 or more legal states. For example, $\{(\{\alpha\}, \{\alpha\}, \alpha)\}$ supports none, $\{(\{\}, \{\}, \alpha)\}$ supports one, and $\{(\{\}, \{\alpha\}, \beta), (\{\}, \{\beta\}, \alpha)\}$ supports two if $\alpha \neq \beta$.

Whenever the reasoner adds or deletes reasons, RMS updates the set of nodes to produce a new legal state. Because a single set of reasons may support any of several sets of nodes, these updates involve choices. The RMS update algorithm was designed to attempt to update states conservatively. That is, if one takes the current state (N, R) and modifies R to obtain R', RMS should choose a new legal state (N', R') with a set of nodes N' as close as possible to the current set N. More precisely, RMS should choose N' so that for any other legal state (N'', R'), neither

[1] It is easy to construct theories in which reasons are nodes themselves, and so may support other reasons. In such theories, one may express all reasons as defeasible reasons and express all changes of belief through addition of reasons. See (Doyle, 1983b) for details.

$N \triangle N'' \subseteq N \triangle N'$ nor $N \triangle N' \subseteq N \triangle N''$ holds, where $\triangle$ denotes symmetric difference ($X \triangle Y = (X \setminus Y) \cup (Y \setminus X)$). Due to the difficulty of quickly computing this form of conservative updates, however, RMS only approximates conservative updates. Reason maintenance systems developed later use simpler notions of conservatism which may be computed more rapidly. For updating beliefs based on monotonic reasons, for example, McAllester's (1982; 1990) efficient so-called "boolean constraint propagation" technique maintains orderings of the premises generating belief sets. As Nebel (1989; 1990) points out, this ordering is analogous to an entrenchment ordering, and permits rapid conservative updates.

The belief revision operations of expansion and contraction correspond to the operations of adding or deleting reasons. Expansion corresponds to adding a reason of the form $(\{\}, \{\}, \alpha)$, which makes α a premise or basic belief. Contraction corresponds to removing all reasons for a node. Addition of other sorts of reasons does not correspond directly to any belief revision operations, as these reasons specify part of the (possibly nonmonotonic) logic of beliefs. Adding a monotonic reason with a nonempty inlist corresponds to adding an ordinary argument step, while adding a nonmonotonic reason lets RMS adopt a belief as an assumption.

3.2 Dependency-directed backtracking

The operation of revision has no direct realization in changes to the set of reasons. It instead corresponds to the operation of the *dependency-directed backtracking* (DDB) system, which we have omitted from the formalization of RMS above.

Unlike the operations of expansion and contraction, the operation of revision involves the notion of logical consistency in an essential way. Consistency does not matter to RMS in adding or deleting reasons; instead, the basic operations of RMS only maintain coherence of reasons and nodes. (Many discussions of reason maintenance misrepresent the truth by claiming that RMS maintains consistency of beliefs. This misrepresentation may stem from the somewhat misleading role of logical consistency in nonmonotonic logic (McDermott and Doyle, 1980).) RMS contraction, moreover, works only for beliefs explicitly represented as nodes, as it depends on removing all explicitly represented reasons for the belief. The operation of revision, in contrast, seeks to resolve a conflict among beliefs rather than to add or subtract reasons from the reason set, and so requires a notion of explicit inconsistency among beliefs.

Since RMS has no knowledge of the meanings of nodes, the reasoner must tell it which nodes represent contradictions. RMS, in turn, informs DDB whenever a contradiction node becomes believed, at which point the backtracker attempts to defeat the arguments supporting the contradiction node by defeating assumptions underlying those arguments. DDB never attempts to remove reasons or premises, only to defeat nonmonotonic assumptions. If the argument for the contradiction node does not depend on any of these (i.e., it consists entirely of monotonic reasons), DDB leaves the contradiction node in place as a continuing belief.

Just as RMS seeks to minimize changes through its incremental update algorithm, DDB seeks to effect minimal revisions in choosing which assumption to defeat. Specifically, the backtracker defeats only "maximal" assumptions that do not depend on other assumptions. This means, in effect, that DDB topologically sorts the beliefs supporting the contradiction node by viewing the reasons comprising the supporting argument as a directed hypergraph. The maximally-positioned assumptions in this ordering of beliefs then constitute the candidates for defeat. The backtracker then chooses one of these assumptions and defeats the reason supporting the chosen assumption by providing a new reason for one of the nodes in its outlist. The new defeating reason is entirely monotonic, and expresses the inconsistency of the chosen assumption with the other maximal assumptions and with other maximally-ordered beliefs in the contradiction node's support. The actual procedure is fairly complex, and we do not attempt to formalize it here. But it is not too misleading to say that the topological sorting process corresponds to calculating an entrenchment ordering, and that DDB revision corresponds to abandoning enough minimally entrenched assumptions to restore consistency. (Cf. (Nebel, 1990, pp. 183-185).)

4 COHERENTIST CRITICISMS OF FOUNDATIONS APPROACHES

Virtually all artificial intelligence approaches to belief revision have been based on some form of reason maintenance, but Harman (1986) and Gärdenfors (1990) reject such foundations-like approaches on a variety of grounds: as psychologically unrealistic (Harman and Gärdenfors), unconservative (Harman), uneconomic (Gärdenfors), and logically superfluous given the notion of epistemic entrenchment (Gärdenfors). We briefly summarize each of these critiques.

4.1 The psychological critique

Experiments have shown that humans only rarely remember the reasons for their beliefs, and that they often retain beliefs even when their original evidential basis is completely destroyed. Harman (1986) interprets these experiments as indicating the people do not keep track of the justification relations among their beliefs, so that they cannot tell when new evidence undermines the basis on which some belief was adopted. Gärdenfors (1990) follows Harman in concluding that the foundations approach cannot be the basis for a psychologically plausible account of belief revision, since it presupposes information that seems experimentally absent.

4.2 The conservativity critique

According to Harman (1986, p. 46), "one is justified in continuing fully to accept something in the absence of a special reason not to", a principle he calls the "principle of conservatism". Harman faults foundations approaches because "the foundations theory rejects any principle of conservatism" ((Harman, 1986, p. 30)), or more precisely, "the coherence theory is conservative in a way the foundations theory is not" ((Harman, 1986, p. 32)). Since Harman views his principle of conservatism as a hallmark of rationality in reasoning, he views the foundations approach as producing irrational behavior.

4.3 The economic critique

According to Gärdenfors (1990), storing derivations of beliefs in memory and using them to update beliefs constitutes a great burden on the reasoner. In the first place, restricting revisions to changing only the foundational beliefs seems too limiting, since nothing in the foundations approach ensures that the foundational beliefs are more basic or valuable in any epistemological or utilitarian sense than the derived beliefs they support. But more importantly, the expected benefit of ensuring that states of belief are well-founded does not justify the cost.

> It is *intellectually extremely costly* to keep track of the sources of beliefs and the benefits are, by far, outweighed by the costs. (Gärdenfors, 1990, p. 31) (his emphasis)
>
> After all, it is not very often that a justification for a belief is actually withdrawn and, as long as we do not introduce new beliefs without justification, the vast majority of our beliefs will hence remain justified. (Gärdenfors, 1990, p. 32)

That is, if a belief has been held and has not caused trouble, there seems to be no good reason to abandon it just because the original reason for which it was believed vanishes or is forgotten. In short, the global extra computations needed to record and ensure proper foundations for belief seems unwarranted compared with a simple approach of purely conservative updates of belief states.

4.4 The superfluity critique

On the face of it, the foundations approach provides reasons for beliefs while the coherence approach does not. For example, in a deductively closed set of beliefs, every belief follows from all the rest, so the coherence approach seems lacking when it comes to providing reasonable explanations of why one believes something. But Gärdenfors (1990) claims that this supposed advantage misconceives the power of the coherence approach, in that the coherence approach can also (in some cases) provide reasons for beliefs.

Gärdenfors observes that any coherence approach involves not only the set of beliefs but also something corresponding to the ordering of epistemic entrenchment, and that while the belief set itself need not determine reasons for beliefs, the entrenchment ordering can be examined to reveal these reasons. For example, let r stand for "it rains today" and h stand for "Oscar wears his hat", and suppose that a deductively closed state of belief K contains r, $r \rightarrow h$, and h. Consider now the result of removing r from K. Superficially, conservatism says that we should retain as many current beliefs as possible, including the belief h. But the result of this contraction depends on how entrenched different beliefs are, and Gärdenfors shows that whether h is retained in $K \dot{-} r$ depends on whether $r \vee \neg h$ is more entrenched than $r \vee h$. If rain is the *only* reason for Oscar to wear his hat, then $r \vee \neg h$ will be more entrenched than $r \vee h$, and we will have $h \notin K \dot{-} r$. If Oscar wears his hat even if it does not rain, then $r \vee h$ will be the more entrenched, and we will have $h \in K \dot{-} r$.

Gärdenfors generalizes from this example to suggest that it may be possible to give a general definition of reasons in terms of epistemic entrenchment. He discusses one possibility for such a definition due to Spohn (1983). Reformulated in the AGM framework, this definition just says that α is a reason for β in K if $\alpha, \beta \in K$ and $\alpha \vee \beta \leq \alpha$; that is, if β is also removed whenever α is. This reconstruction of reasons from epistemic entrenchment does not fully reproduce the foundations approach, however. As Gärdenfors points out, according to this criterion we may have α is a reason for β is a reason for γ is a reason for α, which shows that this particular construction does not correspond to tracing beliefs back to foundational, self-justifying beliefs. Nevertheless, Gärdenfors hopes that a better definition may be possible, and concludes that "representing epistemic states as belief sets together with an ordering of epistemic entrenchment provides us with a way to handle most of what is desired both by the foundations theory and the coherence theory" (Gärdenfors, 1990, p. 45). He thus concludes that the ability to distinguish reasons does not constitute a true advantage of the foundations approach.

4.5 A closer look

While the coherentist critiques certainly give one cause to hesitate before embracing reason maintenance, one may also wonder why artificial intelligence researchers have adopted it so frequently if it has so many defects and no offsetting advantages. To understand better the relative attractions of these approaches, we take a closer look at the issues raised by these critiques. We find, in fact, that these critiques do not stand.

5 ASSESSING THE PSYCHOLOGICAL CRITIQUE

Harman and Gärdenfors correctly observe that the foundations approach, at least as embodied in systems like RMS, corresponds poorly to observed human behavior. But the force of this observation is not at all clear, since psychological accuracy need not be the only aim of a theory of belief revision. In particular, the aim of most artificial intelligence work on belief revision has been to construct computationally useful reasoning mechanisms. One might be pleased if the most useful mechanisms turn out to be psychologically plausible, but computational utility does not depend on that. Indeed, humans might prefer to use a computational system that reasons differently than they do if they find they do better using it than relying on their own abilities.

Moreover, it seems unreasonable to criticize the foundations approach on both psychological and economic grounds. If psychological accuracy is the aim, it matters little if alternative means provide greater efficiency. Efficiency and psychological accuracy bear particularly little correlation when different underlying embodiments (machines or brains) are considered, since humans may not be able to exploit the most efficient computational techniques in their own thinking.

But even if one does take the aim of the theory to be psychological accuracy, recording and using reasons in belief revision does not entail producing unnatural behavior. RMS ensures only that all beliefs enjoy well-founded support, not that all arguments are well-founded. Nothing prevents one from supplying RMS with a circular set of reasons, such as $\{(\{\}, \{\alpha\}, \beta), (\{\}, \{\beta\}, \alpha)\}$. Indeed, most problems yield circular sets of reasons quite naturally simply by reflecting logical relationships among beliefs. As in coherence approaches, which beliefs RMS (or the reasoner) takes as basic may change, at least in principle, with the task at hand.

More fundamentally, however, the recording and using of reasons to revise beliefs does not presuppose a foundations approach. Recent work continues to borrow from coherence approaches by recognizing that recording reasons does not commit one to actually using them, or to making them accessible when reporting on one's reasons. For example, where the original RMS used recorded reasons compulsively in updating beliefs, as prescribed by the foundations approach, the theories of reason maintenance developed in (Doyle, 1983b) make the degree of grounding variable, so that beliefs may be either "locally" or "globally" grounded. Similarly, the rational, distributed reason maintenance service described in (Doyle and Wellman, 1990) uses reasons to revise beliefs only as seems convenient, so that the effective foundations of the belief set change due to further processing even without new information about inferences. This system therefore violates the foundations requirement that beliefs hold only due to well-founded arguments for them. This suggests that one clearly separate the foundations requirement from the more general notions of recording and using reasons as aids to explanation and recomputation, since the latter roles for reasons make sense in both foundations and coherence approaches.

6 ASSESSING THE CONSERVATIVITY CRITIQUE

Addressing the conservativity critique requires recognizing that discussions of belief revision employ two senses of the term "conservatism". Harman uses the term to mean that beliefs *persist* and do not change without specific reasons to disbelieve them. Using this sense of the term, we may agree with his claim that foundations approaches reject conservatism, for the foundations approach calls for abandoning beliefs when one no longer has specific reasons for believing them. But this criticism takes a narrow view of what constitutes a reason for belief. One may take Harman's principle of conservatism as stating a general reason for belief: namely, a lack of past indications that the belief is false constitutes a reason for continuing belief. Such reasons can be respected by a foundations approach simply by adding explicit defeasible reasons for each belief. That is, one might reproduce the coherence approach within the foundations approach simply by adding a reason

$$(\{\}, \{\text{“}\alpha \text{ has been directly challenged”}\}, \alpha)$$

for each node α the first time the α becomes *in*.

Gärdenfors, on the other hand, uses the term "conservatism" to mean *minimal change*, that beliefs do not change more than is necessary. This sense corresponds to the notions of contraction and revision captured by the AGM postulates, according to which one does not abandon both α and β when abandoning α would suffice. This sense of the term seems somewhat more general than the alternative persistence sense, since the persistence interpretation seems to cover only the operation of revision, ignoring the operation of contraction. (Harman, of course, also endorses minimal changes for revision operations, even if his principle of conservatism seems to rule out contraction operations.) But, as the preceding description of RMS makes clear, reason maintenance involves conservative updates in this sense, as its incremental revision approach minimizes the set of changed beliefs. This same behavior can apply to any foundations approach, so that this sense of conservatism does not distinguish coherence and foundations approaches at all.

7 ASSESSING THE ECONOMIC CRITIQUE

Gärdenfors (1990) makes the claim that reason maintenance entails excessive costs, but provides no support for this claim. He describes the reason maintenance system of (Doyle, 1979) as an exemplar of the foundations approach, but presents no analysis of costs of either that specific system or of the foundations approach more generally. More to the point, he makes no attempt to compare the cost of the foundations approach with the cost of the AGM approach he favors. We do this now, and see that the situation is considerably different than that implied by the economic critique.

7.1 Logical costs

Both the coherence and foundations approaches have as their ideal deductively closed states of belief. Unfortunately, practical approaches must abandon this ideal since it entails infinite belief states. Practical approaches must instead be based on finite versions or representations of these theories. Fortunately, restricting attention to finite cases poses no insurmountable problems for either approach. The AGM postulates for contraction, for example, apply perfectly well to finite belief bases, and RMS revises finite belief bases together with finitely many of their conclusions.

Deductive closure, however, is not the only logical property one must abandon to achieve a practical approach, for the requirement that states of belief be logically consistent also entails considerable computational costs for both the coherence and foundations approaches in the usual philosophical conception. Since we model beliefs as sentences in a logical language, determining the consistency or inconsistency of a set of sentences will be undecidable for even moderately expressive languages, and will require time exponential in the size of the sentences in the worst case for sentences of propositional logic. To obtain a realistic theory in accord with the observed ease of human belief revision, it may be necessary to weaken or drop the consistency requirement, as is done in many artificial intelligence approaches. RMS, for example, lacks any knowledge of the what its nodes mean, depends on the reasoner to tell it when some node represents a contradiction, and leaves the conflicting beliefs in place if they do not depend on defeasible assumptions.

Problems remain for the economic critique even if we drop the requirements that beliefs be closed and consistent. In the first place, the AGM approach requires, in postulate ($\leq$2), that logically more general beliefs must also be more entrenched. Postulates ($\leq$3)-($\leq$5) also involve logical dependencies among beliefs. Since determining entailment is as difficult as determining consistency, it would seem that practical approaches must also drop these restrictions on entrenchment orderings. For example, Nebel's (1989; 1990) theory of belief base revision employs orderings of *epistemic relevance*, which are like entrenchment orderings, but which need not respect any logical dependencies among beliefs. Nebel shows that revision according to epistemic relevance orderings satisfies the AGM postulates when lifted to deductively closed states of beliefs. As another example, the theory of *economically rational* belief revision presented in (Doyle, 1991) replaces epistemic entrenchment orderings with arbitrary preference orderings over beliefs and belief states. Like epistemic relevance orderings, these preference orderings need not respect any logical dependencies. But because contraction and revision are defined differently in this theory (as choices rational with respect to the preferences), economically rational belief revision need not satisfy the AGM postulates except when preferences do respect logical dependencies (that is, when one prefers logically more informative states of belief to less informative states).

We conclude that to obtain a practical approach to belief revision, we must give up both logical closure and the consistency and dependency requirements of the AGM approach. If we do not, AGM belief revision is manifestly more costly than more logically modest approaches like reason maintenance.

7.2 Computational costs

To get to the underlying question of whether revising beliefs via epistemic entrenchment costs less computationally than revising beliefs via foundations approaches, we must place the approaches on equal footings and ignore the logical structure of propositions. This means considering arbitrary propositional orderings in the coherence approach, and nonlogical systems like RMS in the foundations approach.

One may easily determine upper bounds on the computational costs of reason maintenance. Updating beliefs after adding or removing a reason costs little when all reasons are monotonic: the time required is at worst cubic in the number of beliefs mentioned by reasons in R, and typically is much less. Updating beliefs apparently costs more when nonmonotonic reasons are used: in the typical system, the time required is at most exponential in the number of beliefs mentioned by reasons in R.

The costs of revising beliefs via epistemic entrenchment are, in contrast, much harder to analyze. To begin with, one must first recall that the AGM approach permits the entrenchment ordering to change with the belief state. To assess the total cost of revision, therefore, one must take the cost of updating this ordering into account. Since the AGM approach never specifies anything about how entrenchment orderings change, we must make assumptions about these changes to conclude anything at all about the cost of entrenchment-based belief revision.

If the ordering of epistemic entrenchment never changes, the cost of belief revision is at worst exponential in the size of the belief base. Gärdenfors and Makinson (1988) show that one can represent the entrenchment ordering in a size linear in the set of dual atoms of a finite belief base. Since these dual atoms are generally of a size exponential in the size of the belief base, examining this representation requires time exponential in the size of the belief base. Comparing this with the cost of reason maintenance, we see that fixed entrenchment revision costs more than reason maintenance with only monotonic reasons, and costs no less than reason maintenance with nonmonotonic reasons.

The economic critique, however, does not seem compatible with the assumption of fixed epistemic entrenchment orderings. Given the emphasis in (Gärdenfors, 1990) on how epistemic entrenchment can represent information about what beliefs justify others, it seems reasonable to assume that the entrenchment ordering must change at least as frequently as does information about justifications. Any application characterized by an ordering that remains constant throughout reasoning is analogous to belief revision given a fixed set of reasons, changing only premises. As just noted, these revisions can be computed as quickly using reason maintenance as when using the dual-atom representation of the entrenchment ordering. Thus to complete the comparison we must determine the cost of updating epistemic entrenchment orderings.

We cannot offer any precise analysis of the cost of updating entrenchment orderings. But if entrenchment orderings do in fact capture all the information in reasons, then updating entrenchment orderings must be roughly as costly as updating reasons. In this case, translating reasons to entrenchment orderings and then using the orderings to compute revisions provides an algorithm for computing revisions from reasons. Any lower bound on the cost of computing revisions from reasons thus provides (ignoring the costs of translation) a lower bound on the cost of computing revisions from entrenchment orderings. Updating entrenchment could cost less only if entrenchment orderings cannot capture all the content of reasons, or if the cost of translating between reasons and orderings is comparable to the cost of revision. Thus if reasons do reduce to entrenchment orderings, then revision using entrenchment may be at least as costly as revision using reasons. This is, of course, just the opposite of the conclusion of the economic critique.

7.3 Practical convenience

Any practical assessment of belief revision and reason maintenance must take into account the human costs of the approach in addition to purely computational costs. One may prefer to use a somewhat more computationally costly approach if it offers much greater convenience to the user. We consider two issues connected with the practical convenience of coherence and foundations approaches.

In the first place, the practical utility of a revision method depends in part on how hard the user must work to give the system new information, in particular, new reasons or new entrenchment orderings. An entrenchment ordering consists of a sequence of sets of equivalently entrenched beliefs, and one might well think this a simpler structure than a set of reasons that this ordering might represent. But if the typical update to this ordering stems from new inferences drawn by the reasoner, then reasons offer the simpler structure for describing the new information, since they correspond directly to the structure of the inference. If the entrenchment ordering is to be updated with every new inference, it seems plausible that the underlying representation should be a set of reasons, from which entrenchment orderings are computed whenever desired. Indeed, we expect that foundations approaches like reason maintenance provide the most natural way of representing and updating entrenchment orderings, particularly when there may be multiple reasons for holding beliefs.

In the second place, the practical utility of a revision method depends in part on how hard the user must work to represent specific information in the system's language. The complete orderings used to describe epistemic entrenchment and epistemic relevance offer limited flexibility in expressing the revision criteria that prove important in practice. Voting schemes, for example, appear often in common methods for choosing among alternatives, but no fixed complete ordering of propositions can express majority voting. In particular, fixed orderings cannot express conservative revision principles like minimizing the *number* of beliefs changed during revision (i.e., Harman's (1986) "simple measure"). Achieving any dependence of ordering on the global composition of the alternatives means revising the linear propositional order to fit each set of alternatives, which hardly seems a practical approach if the user must supply these reorderings. Requiring one to express revision information in the form of entrenchment orderings may therefore impede construction of useful revision systems when the task calls for criteria beyond those easily expressible.

The need for flexibility in specifying contraction and revision guidelines becomes even more apparent if we look to the usual explanations, such as those in (Quine and Ullian, 1978), of why one revision is selected over another. Many different properties of propositions influence whether one proposition is more entrenched than another; one belief might be more entrenched because it is more specific, or was adopted more recently, or has longer standing (was adopted less recently), or has higher probability of being true, or comes from a source of higher authority. As we may expect to discover new guidelines specific to particular tasks and domains, we might expect a practical approach to belief revision to provide some way of specifying these guidelines piecemeal and combining them automatically.

Three problems with these specific guidelines complicate matters, however. The first problem is that these revision criteria are often partial. For example, there are many different dimensions of specificity, and two beliefs may be such that neither is more specific than the other. Similarly, probabilities need not be known for all propositions, and authorities need not address all questions. Thus while each of the guidelines provides information about entrenchment orderings, each provides only partial ordering. In formal terms, each guideline corresponds to a preorder in which not all elements are related to each other.

The second problem is that none of the specific guidelines constitute a comprehensive criterion that takes all possible considerations into account. A comprehensive picture of the entrenchment ordering, all things considered, comes only by combining all of the specific guidelines. Moreover, the overall ordering may be partial, just like the specific guidelines.

The third problem is that the specific guidelines may conflict in some cases. To borrow an example from nonmonotonic logic, one guideline might order beliefs that Quakers are pacifists more entrenched than beliefs that Quakers are not pacifists, and another guideline might order beliefs that Republicans are not pacifists more entrenched than beliefs that Republicans are pacifists. These orderings conflict on cases like that of Nixon, and a guideline ordering more specific rules more entrenched than more general ones does not help since "Quaker" and "Republican" are incomparable categories. Indeed, as argued in (Doyle and Wellman, 1991), other ordering criteria can conflict as well, including very specific criteria corresponding to individual nonmonotonic reasons and default rules. Constructing a global ordering thus means resolving the conflicts among the specific guidelines being combined.

Thus if we seek truly flexible contraction and revision, we need some way of modularly combining and reconciling partial, conflicting, noncomprehensive orderings of propositions into complete global orderings. Unfortunately, it appears unlikely that modular combination and reconciliation can always be done in a rational manner, as Doyle and Wellman (1991) reduce this problem to the problem of social choice, for which Arrow's theorem indicates no good method exists. See also (Doyle, 1991) for a discussion in terms of economically rational belief revision.

To sum up, practical belief revision must depart from the idealizations imposed by epistemic entrenchment. We do not yet know how to specify the necessary information in the most convenient fashion. But preliminary considerations suggest that reasons, not complete entrenchment orderings, offer the most convenient representation.

8 ASSESSING THE SUPERFLUITY CRITIQUE

The superfluity critique says that coherence approaches based on epistemic entrenchment already contains the information needed to identify reasons, obviating the main motivation for foundations approaches. In this section, we identify both some causes to doubt that epistemic entrenchment actually renders foundations approaches superfluous, and reasons to believe that foundations approaches can serve to represent information about epistemic entrenchment.

The program outlined by Gärdenfors for using epistemic entrenchment identify reasons for beliefs faces severe problems. Gärdenfors himself points out the first of these, namely that we do not now possess an adequate definition of reasons in terms of epistemic entrenchment, as the best current candidate does not always distinguish basic from derived beliefs. But this program faces other obstacles as well: an inability to treat multiple reasons for the same belief properly, and an inability to identify temporarily invalid reasons.

To understand the problem of multiple reasons, we consider whether a proposition α is a reason for $\alpha \vee \beta$. Suppose that α and β are epistemically independent propositions. Intuitively, α and β are independent reasons for $\alpha \vee \beta$, or for third propositions γ such that $\alpha \rightarrow \gamma$ and $\beta \rightarrow \gamma$ (so that $\alpha \vee \beta$ is a reason for γ). But neither α nor β need be a reason for $\alpha \vee \beta$ according to the (admittedly flawed) Spohn-Gärdenfors definition, which works only when one belief is the *only* reason for the other. If the two independent propositions are equally entrenched, then neither is a reason for the disjunction; contracting by either leaves the other unaffected, and so also the disjunction, so we have $\alpha \sim \beta < \alpha \vee \beta$, which means $\alpha \vee (\alpha \vee \beta) \not\leq \alpha$ and $\beta \vee (\alpha \vee \beta) \not\leq \beta$. While some better definition of reasons in terms of entrenchment might overcome this difficulty, the prospects seem dim for a definition that involves only the relative entrenchment of the two propositions and their logical combinations, since it appears that adding additional reasons for a belief can cause complex changes in the corresponding entrenchment ordering.

The problem of invalid reasons concerns whether $\alpha \rightarrow \beta$ is a reason for β. In the sense of "reason" formalized in RMS (according to which "reasons" are not themselves beliefs but instead are best viewed as inference rules, constitutive intentions, or constitutive preferences (Doyle, 1988; Doyle and Wellman, 1991)) , the answer is yes, since one may have a reason $(\{\alpha\}, \{\}, \beta)$ independent of whether either α or β is believed. In the sense of "reason" employed in the superfluity critique, the answer is no, since a proposition can have a reason only if both it and the reason are believed. The RMS sense would seem to be the more useful, particularly as a guide to hypothetical reasoning, but one would have to modify the notion of epistemic entrenchment, abandoning postulate ($\leq$4), in order to capture such distinctions.

Even if the superfluity critique is correct in supposing that reasons can be encoded in epistemic entrenchment orderings, the force of the critique is weak unless one also shows that reasons cannot in turn encode entrenchment orderings. But this seems false. We cannot offer a definitive answer here, however, and simply suggest a couple possible avenues towards encoding entrenchment relations in reasons.

The most obvious approach takes the Spohn-Gärdenfors definition at face value and assigns a set of reasons $R_{\leq}$ to each entrenchment ordering $\leq$ over beliefs K such that

$$R_{\leq} \stackrel{\text{def}}{=} \{(\{\alpha\}, \{\}, \beta) \mid \alpha, \beta \in K \ \wedge \ (\alpha \vee \beta \leq \alpha)\}.$$

While this approach might be made to work, doing so requires reconciling the different senses of the term "reason", as noted earlier.

Perhaps a better approach is to simply define the entrenchment ordering in terms of appropriate behaviors of RMS, just as (1) defines the ordering in terms of the results of contraction. For example, we might say that $\alpha \leq \beta$ holds in a RMS state (N, R) just in case removing all reasons for β also removes α, for formally, just in case $\alpha \notin N'$ for every legal state $(N', R \setminus \{(I, O, \beta) \mid I, O \subseteq \mathcal{N}\})$. For example, if

$$R = \left\{ \begin{array}{c} (\{\alpha\}, \{\}, \beta) \\ (\{\beta\}, \{\}, \alpha) \end{array} \right\}$$

then $\alpha \sim \beta$, while if

$$R = \left\{ \begin{array}{c} (\{\alpha\}, \{\}, \beta) \\ (\{\beta\}, \{\}, \alpha) \\ (\{\}, \{\}, \beta) \end{array} \right\}$$

then $\alpha < \beta$, and if

$$R = \left\{ \begin{array}{c} (\{\alpha\}, \{\}, \beta) \\ (\{\beta\}, \{\}, \alpha) \\ (\{\}, \{\}, \alpha) \\ (\{\}, \{\}, \beta) \end{array} \right\}$$

then $\alpha \sim \beta$ once again. If one also desires to capture the logical structure of epistemic entrenchment, one can also augment the set of nonlogical reasons with a set of monotonic reasons describing all logical dependencies among the propositions. That is, one expands the set of nodes to include new nodes representing all logical equivalence classes elements of the propositional algebra over the original nodes, and expands the set of reasons to include a reason $(I, \{\}, \alpha)$ whenever $I \vdash \alpha$ in the expanded set of nodes.

To sum up, a satisfactory reduction of reasons to entrenchment orderings remains to be provided, and this will support the superfluity critique only if one cannot similarly reduce entrenchment orderings to reasons.

9 CONCLUSION

Both coherence and foundations approaches to belief revision provide valuable perspectives. The AGM approach focuses on the ordering of beliefs according to epistemic entrenchment, while reason maintenance focuses on the reasons relating individual beliefs. Though reason maintenance has been criticized on grounds of cost, psychological accuracy, and logical necessity, a close examination of these criticisms reveals the situation to be much more complex than that portrayed by the criticisms. While a definitive conclusion about whether either of these approaches is better than the other awaits answers to questions about whether either one mathematically subsumes the other, we believe that reason maintenance incorporates the key elements of the coherence approach, while at the same time providing the most practical means of mechanizing coherence approaches.

Perhaps the most fruitful way of viewing the issue is to focus on the great and fundamental similarities of the approaches rather than on their apparently minor differences. At least as far as the specific AGM and RMS approaches are concerned, we see that:

- Both seek to recognize logical and inferential relations among beliefs, differing at most in how these relations are represented. Both must abandon most requirements of logical consistency and closure to be useful in practical mechanizations.
- Both make minimal changes of belief, differing at most in the set of possible alternatives entering into the minimization.
- Both allow flexibility in choosing whether to reflect reasons and other inferential relations in epistemic states, differing only in whether representations of reasons determine entrenchment orderings or representations of entrenchment orderings determine reasons. Both make no stipulations about what reasons or entrench-

ment orderings should be represented, other than to assume this information may change with each step of reasoning.

- Neither distinguishes in any fixed way between "fundamental" and "derived" beliefs. Instead, both allow one to ground beliefs only to the extent required by one's needs for explanations and updates, and to change the identification of basic beliefs along with the current purposes of reasoning.

Given these similarities, the important questions for artificial intelligence concern the relative computational efficiencies of different representational schemes: not just AGM coherence and traditional reason maintenance, but possibly mixed schemes as well. Getting a clearer theoretical picture of such schemes and their relative merits promises to yield many valuable practical returns.

Acknowledgments
I thank Peter Gärdenfors, David Makinson, Bernhard Nebel, Robert Stalnaker, Richmond Thomason, and Michael Wellman for valuable discussions of this topic. This work was supported by the USAF Rome Laboratory and DARPA under contract F30602-91-C-0018, and by the National Library of Medicine through National Institutes of Health Grant No. R01 LM04493.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction functions and their associated revision functions. *Journal of Symbolic Logic* 50:510–530.

Doyle, Jon 1979. A truth maintenance system. *Artificial Intelligence* 12(2):231–272.

Doyle, Jon 1983a. The ins and outs of reason maintenance. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*. 349–351.

Doyle, Jon 1983b. Some theories of reasoned assumptions: an essay in rational psychology. Technical Report 83-125, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

Doyle, Jon 1988. Artificial intelligence and rational self-government. Technical Report CS-88-124, Carnegie-Mellon University Computer Science Department.

Doyle, Jon 1991. Rational belief revision (preliminary report). In Fikes, Richard E. and Sandewall, Erik, editors, *Proceedings of the Second Conference on Principles of Knowledge Representation and Reasoning*, San Mateo, CA. Morgan Kaufmann. 163–174.

Doyle, Jon and Wellman, Michael P. 1990. Rational distributed reason maintenance for planning and replanning of large-scale activities. In Sycara, Katia, editor, *Proceedings of the DARPA Workshop on Planning and Scheduling*, San Mateo, CA. Morgan Kaufmann. 28–36.

Doyle, Jon and Wellman, Michael P. 1991. Impediments to universal preference-based default theories. *Artificial Intelligence* 49(1–3):97–128.

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

Gärdenfors, Peter 1990. The dynamics of belief systems: Foundations vs. coherence theories. *Revue Internationale de Philosophie* 172:24–46.

Gärdenfors, Peter and Makinson, David 1988. Revisions of knowledge systems using epistemic entrenchment. In Vardi, Moshe Y., editor, *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, Los Altos, CA. Morgan Kaufmann. 83–95.

Harman, Gilbert 1986. *Change in View: Principles of Reasoning.* MIT Press, Cambridge, MA.

McAllester, David 1982. Reasoning Utility Package user's manual. Artificial Intelligence Memo 667, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA.

McAllester, David 1990. Truth maintenance. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, Menlo Park, CA. AAAI Press. 1109–1116.

McDermott, Drew and Doyle, Jon 1980. Non-monotonic logic—I. *Artificial Intelligence* 13:41–72.

Nebel, Bernhard 1989. A knowledge level analysis of belief revision. In Brachman, Ronald J.; Levesque, Hector J.; and Reiter, Raymond, editors, *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, San Mateo, CA. Morgan Kaufmann. 301–311.

Nebel, Bernhard 1990. *Representation and Reasoning in Hybrid Representation Systems.* Number 422 in Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin.

Quine, Willard V. 1953. Two dogmas of empiricism. In *From a Logical Point of View: Logico-Philosophical Essays.* Harper and Row, New York, second edition. 20–46.

Quine, W. V. 1970. *Philosophy of Logic.* Prentice-Hall, Englewood Cliffs, NJ.

Quine, W. V. and Ullian, J. S. 1978. *The Web of Belief*. Random House, New York, second edition.

Spohn, W. 1983. Deterministic and probabilistic reasons and causes. *Erkenntnis* 19:371–396.