Introduction to Belief Revision*

Maurice Pagnucco

Department of Computing Science
Division of Information and Communication Sciences
Macquarie University
NSW 2109

E-mail: morri@ics.mq.edu.au
WWW: http://www.comp.mq.edu.au/∼morri/

> *You could not step twice into the same rivers; for other waters are ever flowing on to you.*
> Heracleitus of Ephesus (as quoted by Plutarch, *On the E at Delphi*)

The concept of belief revision is, in essence, a simple one. We are interested in characterising the dynamics of epistemic states; how an agent in a particular epistemic state modifies this state upon receipt of some new information (or epistemic input). Moreover, we are interested in investigating changes of belief that are performed in a *rational* manner.

1 Foundationalism Versus Coherentism

Before addressing the problem of how to alter an epistemic state given an epistemic input, we shall briefly investigate the nature of the states themselves. The two foremost approaches to modelling epistemic states are the *foundational* and *coherence* theories. Pollock [46] refers to these as *doxastic* theories; they assume that the justificatory pedigree of beliefs depends solely on those beliefs held by an agent.

---

[*] This paper is taken, with minor alterations, from Chapter 2 of Maurice Pagnucco, *The Role of Abductive Reasoning Within the Process of Belief Revision*, PhD thesis, Department of Computer Science, University of Sydney, February 1996.

The major distinguishing feature of the foundational approach is that it demarcates a special class of beliefs. These are often referred to as "epistemologically basic beliefs" (or simply "basic beliefs"). Every belief in a foundational system is supposed to be justified in terms of other beliefs which are, in turn, justified by further beliefs until ultimately we reach basic beliefs which have no need of justification. In a certain sense, they can be thought of as self justifying.

Examples of foundational systems within artificial intelligence are Doyle's [7] Truth Maintenance System (TMS)[1] and its successor the Assumption Based Truth Maintenance System (ATMS) [49]. Their basic task is to record inferences passed to them by a domain dependent problem solver. Basically, the TMS consists of two structures: *nodes* representing propositions; and, *justifications* representing reasons. Each node may be in one of two states:

in	the node has a valid justification and consequently is considered a current belief
out	the node does not have a valid justification; it is currently not believed

A justification is a pair of sets of nodes: an *inlist* and an *outlist*. A justification is *valid* if all nodes on its *inlist* are *in* and all those on its *outlist* are *out*. Clearly, a proposition becomes a belief when one of its justifications is valid; it becomes a non-belief when none of its justifications are valid. The TMS takes care of creating new nodes and adding or retracting justifications. This may become a complex process as other nodes and justifications may be affected. Also, circular justifications must not be admitted. It can also mark a node as a *contradiction*. This has the effect of stating that the elements of a justification for this node are inconsistent. If such a node acquires a valid justification a process known as *dependency-directed backtracking* ensues, making sure that any justification is no longer valid.[2]

Elkan [9] (see also Reinfrank *et al.* [47, 48]) provides a logical rendering of the TMS in order to show how it relates to Gelfond and Lifschitz' [17] stable model semantics for logic programming and Moore's [36] (propositional) autoepistemic logic. A justification for a proposition c is simply represented as a (directed) propositional clause

$$a_1 \wedge \ldots \wedge a_n \wedge \neg b_1 \wedge \ldots \wedge \neg b_m \rightarrow c$$

(where $\neg$ represents negation as failure).[3] Propositions $a_1, \ldots, a_n$ represent those

---

[1] The term Truth Maintenance System has often been cited as a misnomer and the alternative Reason Maintenance System (RMS) suggested as a more appropriate alternative. However, the term Truth Maintenance System appears to have stuck and we shall use it here.

[2] This is essentially achieved by retracting beliefs known as *assumptions* — having a justification with a non-empty *outlist* — which, although not necessarily part of the justification, lead to it becoming invalid.

[3] In autoepistemic logic this justification may be rendered $a_1 \wedge \ldots \wedge a_n \wedge \neg \mathrm{L} b_1 \wedge \ldots \wedge \neg \mathrm{L} b_m \rightarrow c$.

in the justification's *inlist* while $b_1, \ldots, b_m$ are those in the *outlist*. A contradiction c can be represented by the clause $a_1 \wedge \ldots \wedge a_n \rightarrow c$ meaning that $a_1, \ldots, a_n$ cannot all be believed (think of a contradiction c as representing $\perp$).

One of the more popular successors of the TMS is de Kleer's [4] ATMS. It is based on the idea of keeping track of the assumptions upon which a proposition is based as well as its justifications.[4] In this case a node is composed of three parts: the *datum* representing a proposition but treated as atomic; the *label* representing sets of assumptions — called *environments* — which would allow the datum to be inferred; and, the *justifications*, each containing antecedents supporting the datum. An *assumption* is a node with an environment consisting of its datum only.[5] The nodes derivable from an environment (including those corresponding to the elements of the environment itself) are referred to as *contexts*. Sets of assumptions that cannot hold simultaneously are referred to as *nogoods*. They are similar to contradictions in the TMS and act like integrity constraints, reducing the size of the search space for any eventual query passed to the ATMS (by causing the deletion of derived justifications that violate these constraints). When an inference is passed to the ATMS it takes care of updating nodes. If the consequent is unknown, a new node is created with the consequent as the datum. The antecedent of the inference becomes a new justification for the node and labels of the antecedents are used in determining the label for the node. All combinations of an environment from every label are used in determining new environments. However, a label must be

consistent	no environment is a superset of a nogood
complete	the environment from which the datum follows is a superset of some environment in the datum's label
sound	the datum follows from each environment
minimal	no environment is a subset of another environment in the label

Soundness and completeness are guaranteed by the procedure used to compute environments. Inconsistent and non-minimal environments must be removed.

Logically, ATMS justifications are simply Horn clauses

$$a_1 \wedge \ldots \wedge a_n \rightarrow c$$

where c represents the datum and $a_1, \ldots, a_n$ a justification. Assumptions can be represented in the same manner however, in this case, $a_1, \ldots, a_n$ is taken from one of the datum's environments. Nogoods $\{a_1, \ldots, a_n\}$ are just supports for falsity, $a_1 \wedge \ldots \wedge a_n \rightarrow \perp$.

---

[4] Moreover, whereas the TMS concentrates on finding one support, the ATMS is geared towards finding all supports.

[5] Cf. TMS — the notions differ.

Example 1.1 *Using part of our previous hamburger example*

$$o_1 \wedge ph_1 \to e$$
$$o_2 \wedge ph_2 \to e$$

Suppose also that the two restaurants serve milkshakes so that, if a restaurant is open and I purchase a milkshake from it (pm_1/pm_2), then I will have a milkshake (m)

$$o_1 \wedge pm_1 \to m$$
$$o_2 \wedge pm_2 \to m$$

Moreover, suppose there is a municipal restriction stating that the two restaurants can never be open at the same time. We express this fact through the addition of the following nogood

$$o_1 \wedge o_2 \to \bot$$

Now, suppose a new inference is passed to the ATMS: I am satiated (s) after having a hamburger and a milkshake

$$e \wedge m \to s$$

All possible justifications for e and m produce the following justifications for s

$$\begin{aligned} o_1 \wedge ph_1 \wedge pm_1 &\to s \\ o_1 \wedge ph_1 \wedge o_2 \wedge pm_2 &\to s \\ o_1 \wedge pm_1 \wedge o_2 \wedge ph_2 &\to s \\ o_2 \wedge ph_2 \wedge pm_2 &\to s \end{aligned}$$

However, the second and third violate the nogood and must be removed. All justifications are minimal. □

The coherence approach, unlike the foundational approach underlying the TMS and ATMS, denies the existence of any select set of basic beliefs. On this account, beliefs are justified by the way they interact or "cohere" with other beliefs. In other words, it is the relationship with other beliefs that is important when determining whether a belief is justified.

Pollock [46] distinguishes four types of coherence theories into two groups:

1a) Positive Coherence

> The agent must possess reasons for maintaining a belief. That is, each belief must have "positive support".

1b) Negative Coherence

> The agent is justified in holding a belief provided there is no reason to think otherwise. ("All beliefs are 'innocent until proven guilty"', Pollock [46] p. 72.)

2a) Linear Coherence

> The agent adopts a more traditional (i.e., foundational) view of reasons except that if we look at a reason, the reasons for holding reasons, etc., we would never stop; either we have an infinite sequence of reasons or there is some circularity in the reason structure.

2b) Holistic Coherence

> The agent is justified in holding a belief due to some relationship between the belief and all other beliefs held.

It is possible to have coherence theories which possess more than one aspect from this list.

The distinction between the foundational and coherence approaches is often illustrated through two metaphors: the foundationalist "pyramid" and the coherentist "raft". These are succinctly expressed by Sosa [51]:

> For the foundationalist every piece of knowledge stands at the apex of a pyramid that rests on stable and secure foundations whose stability and security does not derive from the upper stories or sections. For the coherentist a body of knowledge is a free-floating raft every plank of which helps directly or indirectly to keep all the others in place, and no plank of which would retain its status with no help from the others.

The latter derives from a metaphor by Neurath [41] used to express the fact that it is not possible (nor desirable) to start from scratch in developing a language for scientific discourse:

> We are like sailors who have to rebuild their ship on the open sea, without ever being able to dismantle it in dry-dock and reconstruct it from the best components.

Pollock [46] notes that this metaphor is more in keeping with the negative coherence view.

2 The AGM Framework for Belief Revision[6]

We shall base our study on the AGM framework for belief revision [1, 12, 13]. This approach is claimed, by Gärdenfors [13] (also [12] p. 35), to be coherentist in nature. The main reason for adopting this approach is that it is a well developed formal framework. Moreover, links have been investigated between belief revision and other areas of artificial intelligence (e.g., nonmonotonic reasoning [15, 35]).

It was mentioned earlier that we are interested in accounts of rational belief change. Gärdenfors and Rott [16] adopt the following rationality criteria or "integrity constraints":

- Where possible, epistemic states should remain consistent
- Any sentence logically entailed by beliefs in an epistemic state should be included in the epistemic state
- When changing epistemic states, loss of information should be kept to a minimum[7]
- Beliefs held in higher regard should be retained in favour of those held in lower regard

The third criterion can be thought of as a manifestation of Occam's razor as applied to the removal of information (rather than the making of hypotheses), and is held in high regard in the AGM framework. In fact, it is often mentioned in connection with this framework. We shall see that a variant of it also applies to the acquisition of new information.

The second criterion leads to the following conception of an epistemic state within the AGM framework. Epistemic states are closed under logical consequence (Cn) and are referred to as *belief sets*.[8] The set of all belief sets is denoted $\mathcal{K}$. One special type of belief set is the absurd belief set $K_\perp$ which contains all formulae in $\mathcal{L}$. This rather idealistic modelling of epistemic states may best be viewed as the agent's *doxastic commitment to full recognition* of the truth of the deductive consequences of what it believes (see Levi [31] p. 8). A lot of attention has also

---

[6] The results in this section may be found in the AGM literature; in particular, they are collected together in Gärdenfors [12] unless otherwise stated.

[7] This is also referred to as the Principle of Informational Economy [12] and, when informational loss is measured by set inclusion, the Principle of Conservation [25]. They are special cases of the Principle of Minimal Change [25] which states that minimal change should occur when beliefs are added as well as removed.

[8] That is, belief sets are simply theories albeit with a special interpretation in mind.

been paid to the study of *belief bases* [11, 21, 37, 40]; sets of formulae that are not necessarily closed under the logical consequence operation.

Given any consistent belief set K, there are three types of *epistemic attitude* toward a sentence α:

(*i*) α is *accepted* (or believed) if $\alpha \in K$

(*ii*) α is *rejected* (or not believed) if $\neg\alpha \in K$

(*iii*) α is *indetermined* if $\alpha \notin K$ and $\neg\alpha \notin K$

Epistemic inputs are represented by a single sentence from the object language. More complex representations may be found in the literature (e.g., [23, 52]).

Belief change operators can be seen as prescribing how a given epistemic state is to be altered given an epistemic input. The AGM considers three types of belief change operators given a belief state K, representing the agent's current epistemic state, and epistemic input α:

Belief Expansion (K^+_α) Incorporation of new belief α into K without retraction of any existing beliefs

Belief Contraction (K^-_α) Removal of belief α from K without introduction of any new beliefs

Belief Revision (K^*_α) Incorporation of new belief α into K with possible removal of existing beliefs in order to maintain consistency

A belief change operator is essentially a function taking a belief set K and epistemic input α to a new belief set K'_α $(+, -, * : \mathcal{K} \times \mathcal{L} \to \mathcal{K})$.[9] These belief change operators are investigated in a number of ways: through rationality postulates and through a variety of constructions. The postulates are then related to the constructions via *representation theorems*. The idea is to study all possible belief change functions — that is, all possible ways of expanding, contracting and revising K by α — in accord with the rationality constraints imposed by the postulates.

2.1 Postulates

Rationality postulates specify constraints that the respective operators should satisfy. They are guided by the rationality criteria outlined above which we adopt in this dissertation as the standards for characterising a rational agent.

---

[9] The restriction that the nature of an epistemic state be the same before and after undergoing change is referred to as the *Principle of Categorical Matching* [16].

2.1.1 Expansion

The expansion of a belief set K by an epistemic input α is denoted K^+_α. Expansion is generally recommended when α is consistent with K. An AGM expansion operator $+ : \mathcal{K} \times \mathcal{L} \to \mathcal{K}$ satisfies the following rationality postulates:

(K$^+1$) For any sentence α and any belief set K, K^+_α is a belief set (closure)
(K$^+2$) $\alpha \in K^+_\alpha$ (success)
(K$^+3$) $K \subseteq K^+_\alpha$ (inclusion)
(K$^+4$) If $\alpha \in K$, then $K^+_\alpha = K$ (vacuity)
(K$^+5$) If $K \subseteq H$, then $K^+_\alpha \subseteq H^+_\alpha$ (monotonicity)
(K$^+6$) For all belief sets K and all sentences α, K^+_α is the smallest belief set that satisfies (K$^+1$) — (K$^+5$) (minimality)

The postulate of closure expresses the fact that $+$ is a function taking a belief set and a sentence as input and produces a belief set. Success states that the epistemic input is accepted in the expanded epistemic state. Inclusion says that no beliefs are retracted and is a form of the Principle of Minimal Change as phrased above. Vacuity represents a boundary case and states that nothing need be done if the epistemic input is already accepted.[10] Monotonicity says that, if one belief state contains at least the same information as another, then its expansion will contain at least the information of the expansion of the other with respect to the same epistemic input. The postulate of minimality can be considered an expression of the Principle of Minimal Change applied to the addition of new beliefs to an epistemic state K; the smallest possible change to accommodate the new information is made. The term “smallest” is understood with respect to set inclusion (of the original epistemic state relative to the expanded epistemic state). This leads to the following representation theorem.

Theorem 2.1 *The expansion function* $+$ *satisfies* (K^+1) — (K^+6) *if and only if* $K^+_\alpha = Cn(K \cup \{\alpha\})$.

Therefore, to calculate an AGM expansion, one need only take the deductive closure of the initial epistemic state and the new information.

2.1.2 Contraction

The contraction of a belief set K by epistemic input α is denoted K^-_α. Contraction is recommended when doubt is raised about a current belief or the agent wishes

---

[10] Postulate (K$^+4$) is superfluous as it follows from postulates (K$^+1$) — (K$^+3$), (K$^+5$) and (K$^+6$).

to temporarily suspend belief in a proposition. It can be used together with expansion to perform revision, as we shall see, and satisfies the following rationality postulates:

(K$^-$1) For any sentence α and any belief set K,
K^-_α is a belief set (closure)
(K$^-$2) $K^-_\alpha \subseteq K$ (inclusion)
(K$^-$3) If $\alpha \notin K$, then $K^-_\alpha = K$ (vacuity)
(K$^-$4) If $\not\vdash \alpha$ then $\alpha \notin K^-_\alpha$ (success)
(K$^-$5) If $\alpha \in K$, $K \subseteq (K^-_\alpha)^+_\alpha$ (recovery)
(K$^-$6) If $\vdash \alpha \leftrightarrow \beta$, then $K^-_\alpha = K^-_\beta$ (extensionality)
(K$^-$7) $K^-_\alpha \cap K^-_\beta \subseteq K^-_{\alpha\wedge\beta}$ (intersection)[11]
(K$^-$8) If $\alpha \notin K^-_{\alpha\wedge\beta}$, then $K^-_{\alpha\wedge\beta} \subseteq K^-_\alpha$ (conjunction)[12]

Closure states that a contraction operation takes pairs of belief sets and formulae to belief sets ($- : \mathcal{K} \times \mathcal{L} \to \mathcal{K}$). Inclusion says that no new beliefs should be introduced into the contracted epistemic state. Vacuity expresses the fact that nothing need be done if the epistemic input is not currently accepted. It is a manifestation of the Principle of Minimal Change. Success states that if it is possible to remove the epistemic input, it will be retracted from the current epistemic state. The only situation in which it is not possible to do so occurs when the epistemic input is a logical truth for, by the second of our rationality criteria above, it will be included in all possible epistemic states. Recovery says that if we were to retract a belief from K and then expand the result by the same formula, all original beliefs would be included in the final epistemic state. This behaviour is also due to the Principle of Minimal Change to a certain extent since this principle dictates that beliefs not be unnecessarily discarded when determining K^-_α. Recovery is arguably the most controversial of the AGM rationality postulates and there are a number of contributions discussing its removal [20, 34]. Makinson [34] refers to a contraction operation satisfying postulates (K$^-$1) — (K$^-$4) and (K$^-$6) as a *withdrawal*. Extensionality expresses the Principle of Irrelevance of Syntax; it is the content rather than the syntactic formulation of the epistemic input that is important in belief change. These first six postulates are often referred to as the *basic postulates for contraction over K*. The remaining two postulates are supplementary postulates. They are best motivated in the style of Nayak [37] (p. 506). Intersection states that, if one does not give up belief in γ when giving up belief in α nor in giving up belief in β, then one should not give up belief in γ when giving up belief

---

[11] Also referred to as *conjunctive overlap* [22].
[12] Also referred to as *conjunctive inclusion* [22].

in the conjunction $\alpha \wedge \beta$. Conjunction states that, if one were to give up α when giving up the conjunction $\alpha \wedge \beta$, then whatever one gives up in giving up α, should also be given up in giving up $\alpha \wedge \beta$.

2.1.3 Revision

The revision of a belief set K by epistemic input α is denoted K^*_α. Revision is particularly important when α is inconsistent with K and the agent wishes to incorporate it in such a way as to end up in a consistent epistemic state. It satisfies the following rationality postulates:

(K*1) For any sentence α and any belief set K,
K^*_α is a belief set (closure)
(K*2) $\alpha \in K^*_\alpha$ (success)
(K*3) $K^*_\alpha \subseteq K^+_\alpha$ (inclusion)
(K*4) If $\neg\alpha \notin K$, then $K^+_\alpha \subseteq K^*_\alpha$ (preservation)
(K*5) $K^*_\alpha = K_\perp$ if and only if $\vdash \neg\alpha$ (vacuity)
(K*6) If $\vdash \alpha \leftrightarrow \beta$, then $K^*_\alpha = K^*_\beta$ (extensionality)
(K*7) $K^*_{\alpha\wedge\beta} \subseteq (K^*_\alpha)^+_\beta$ (superexpansion)
(K*8) If $\neg\beta \notin K^*_\alpha$, then $(K^*_\alpha)^+_\beta \subseteq K^*_{\alpha\wedge\beta}$ (subexpansion)

(K*1) is the familiar postulate of closure ($* : \mathcal{K} \times \mathcal{L} \to \mathcal{K}$). Success states that the new information should be included in the revised epistemic state. Inclusion says that expansion represents an "upper bound" when incorporating new beliefs (this will trivially hold in case the negation of the epistemic input is already accepted). Preservation expresses that, when the negation of the epistemic input is not accepted, revision reduces to expansion. It is the conditional converse of inclusion. Vacuity tells us that the only situation in which revision would end up in the inconsistent epistemic state occurs when the agent is asked to accept logically contradictory information. Extensionality, like its contraction counterpart, is an expression of the Principle of Irrelevance of Syntax. Again, (K*1) — (K*6) are referred to as the *basic postulates for belief revision over K*. Two supplementary postulates for revision exist and can be thought of as generalisations of Inclusion and Preservation. Superexpansion states that any belief included in the revision of K by $\alpha \wedge \beta$ should also be included if we first revise K by α and then expand the result by β. Subexpansion says that, if β is not rejected in revising K by α, then any belief included by first revising K by α and expanding the result by β should also be included in the revision of K by $\alpha \wedge \beta$. That is, $K^*_{\alpha\wedge\beta}$ and $(K^*_\alpha)^+_\beta$ consist of the same beliefs in this case. This postulate is the conditional converse of Superexpansion.

Interestingly enough, not all of the above operators are essential; some may be defined in terms of the other operators.[13] A revision operator, for instance, may be determined from a contraction operator and an expansion operator via the *Levi Identity*:

$$(\mathrm{Def}\ *) \qquad\qquad K^*_\alpha = (K^-_{\neg\alpha})^+_\alpha$$

It states that a revision of K by α can be performed by first removing $\neg\alpha$ (to avoid inconsistency) and incorporating α into the result. The following theorem gives credence to this definition (and Levi's claims).

Theorem 2.2 *Let* $-$ *be a contraction function satisfying postulates* (K^-1) — (K^-4) *and* (K^-6) *and* $+$ *an expansion function satisfying postulates* (K^+1) — (K^+6)*. Then the revision function* $*$ *obtained from* $(\mathrm{Def}\ *)$ *satisfies* (K^*1) — (K^*6)*. Moreover, if* $-$ *satisfies* (K^-7)*, then* $*$ *satisfies* (K^*7) *and if* $-$ *satisfies* (K^-8)*, then* $*$ *satisfies* (K^*8)*.*

Notice that recovery is not required in Theorem 2.2. This tells us that revision operators defined, via the Levi Identity, from AGM contractions and those from withdrawal operators are *revision equivalent*. That is, they determine the same class of revision operators.

Alternatively, it is possible to define a contraction operator using a revision operator and set intersection.[14] This may be achieved by the *Harper Identity*:

$$(\mathrm{Def}\ -) \qquad\qquad K^-_\alpha = K\ \cap\ K^*_{\neg\alpha}$$

which states that contracting K by α consists of those beliefs in K that are retained in revising K by $\neg\alpha$. The motivation for this definition stems from the fact that $K^*_{\neg\alpha}$ represents a minimal change of K required to incorporate $\neg\alpha$ (in a consistent manner) and should therefore include a large part of K that does not entail α.

Theorem 2.3 *Let* $*$ *be a revision function satisfying postulates* (K^*1) — (K^*6)*. Then the contraction function obtained from* $(\mathrm{Def}\ -)$ *satisfies* (K^-1) — (K^-6)*. Moreover, if* $*$ *satisfies* (K^*7)*, then* $-$ *satisfies* (K^-7) *and if* $*$ *satisfies* (K^*8)*, then* $-$ *satisfies* (K^-8)*.*

---

[13] In fact, Levi [30] claims that the only "legitimate" forms of changing an epistemic state are expansion and contraction, a view to which we subscribe. He refers to this as the *commensurability thesis* [31] (p. 65). Such a view places less emphasis on the revision operator which is deemed achievable through a sequence of expansions and contractions. As a result, we place a greater emphasis on these latter two operators in this dissertation.

[14] Recall that the intersection of two belief sets is also a belief set.

2.2 Constructions

Having outlined conditions that the various belief change operators should satisfy, it is interesting to study how operators satisfying these postulates could be constructed. The AGM framework possesses four main constructions: selection functions over maximal subsets of K failing to imply α, Grove's system of spheres, epistemic entrenchment and safe contraction.[15]

2.2.1 Selection Functions

Given the Levi Identity and Theorem 2.1 regarding belief expansion, it is sufficient to concentrate on contraction. One approach to constructing a contraction of belief set K by epistemic input α is to seriously consider the Principle of Minimal Change and look at subsets of K which are as big as possible without entailing α. Such a set can be defined as follows:

Definition 2.1 *A belief set K' is a maximal subset of K that fails to imply α if and only if*

(i) $K' \subseteq K$

(ii) $\alpha \notin K'$

(iii) for any $\beta \in \mathcal{L}$, if $\beta \in K$ and $\beta \notin K'$, then $\beta \rightarrow \alpha \in K'$

The set of all belief sets that are maximal subsets of K failing to imply α are denoted $K \perp \alpha$.

Generally, $K \perp \alpha$ contains more than one maximal subset. The first idea in constructing a contraction function is to apply a *selection function* γ to select one element from $K \perp \alpha$.[16] Intuitively, $\gamma(K \perp \alpha)$ returns the "best" element from $K \perp \alpha$ and is known as a *maxichoice selection function*. The contraction of K by α can be defined as follows

$$\text{(Def Max)} \qquad K^{-}_{\alpha} = \begin{cases} \gamma(K \perp \alpha) & \text{whenever } K \perp \alpha \text{ is nonempty}^{17} \\ K & \text{otherwise} \end{cases}$$

and is referred to as a *maxichoice contraction function over K*. Sure enough, such a function satisfies the basic postulates for belief contraction over K.

---

[15] A construction in terms of nice preorders over models (see [28, 15]) is also presented by Peppas and Williams [45] but we shall not consider it here.

[16] A selection function applied to a set X returns an element of the co-domain whenever X is nonempty.

[17] Note that $K \perp \alpha = \emptyset$ only when $\vdash \alpha$.

Lemma 2.4 *Let K be a belief set. If $-$ is a maxichoice contraction function over K, then it satisfies postulates* (K^-1) — (K^-6) *for belief contraction over K.*

Unfortunately, we obtain the following undesirable results.

Lemma 2.5 *Let K be a belief set and $\alpha \in \mathcal{L}$. If $\alpha \in K$ and K^-_α is defined by means of a maxichoice contraction function, then for any proposition β either $\alpha \vee \beta \in K^-_\alpha$ or $\alpha \vee \neg\beta \in K^-_\alpha$.*

Corollary 2.6 *Let $-$ be a maxichoice contraction function over K. If a revision function $*$ is defined from $-$ by the Levi Identity, then, for any α such that $\neg\alpha \in K$, K^*_α is a complete theory.*

They suggest that maxichoice contractions retain too much information. In the resulting revision, the agent has opinions as to the truth or falsity of every proposition.

It seems natural then to consider a selection function at the other extreme; one returning all elements of $K\bot\alpha$. This is known as a *full meet selection function* and leads to a *full meet contraction function over K* which may be defined as follows.

$$(\text{Def Meet}) \qquad K^-_\alpha = \begin{cases} \bigcap(K\bot\alpha) & \text{whenever } K\bot\alpha \text{ is nonempty} \\ K & \text{otherwise} \end{cases}$$

A full meet contraction function also satisfies the basic postulates for contraction.

Lemma 2.7 *Let K be a belief set. If $-$ is a full meet contraction function over K, then it satisfies postulates* (K^-1) — (K^-6) *for belief contraction over K.*

However, we again have undesirable results.

Lemma 2.8 *Let K be a belief set and $\alpha \in \mathcal{L}$. If $\alpha \in K$ and K^-_α is defined by means of a full meet contraction function, then for any proposition β, $\beta \in K^-_\alpha$ if and only if $\beta \in K$ and $\neg\alpha \vdash \beta$.*

Corollary 2.9 *Let $-$ be a full meet contraction function over K. If a revision function $*$ is defined from $-$ by* (Def $*$)*, then for any α such that $\neg\alpha \in K$, $K^*_\alpha = Cn(\alpha)$.*

In a sense, too much information is removed. This is somewhat at odds with the Principle of Minimal Change.

A remedy lies in making a compromise between these two extremes. We adopt a selection function γ that returns a subset of $K\bot\alpha$. We can think of γ as returning the set of “best” elements of $K \bot \alpha$.[18] This is referred to as a *partial meet selection function*. The resulting contraction — a *partial meet contraction function over K* — may be defined as follows

$$\text{(Def Part)} \qquad K^-_\alpha = \begin{cases} \bigcap \gamma(K \bot \alpha) & \text{whenever } K \bot \alpha \text{ is nonempty} \\ K & \text{otherwise} \end{cases}$$

The following representation theorem says that such functions exactly coincide with the basic postulates for contraction.

Theorem 2.10 *Let K be a belief set and $-$ be a contraction function. Then $-$ is a partial meet contraction function over K if and only if it satisfies postulates* (K^-1) — (K^-6) *for contraction over K.*

It is interesting to further investigate the nature of the selection function γ and how it decides which elements of $K \bot \alpha$ are preferred. One idea is to impose a relation $\preceq$ over the elements of $K \bot \alpha$ and define γ by the following *marking-off identity* (when $K \bot \alpha \neq \emptyset$):

$$\text{(Def } \gamma) \qquad \gamma(K \bot \alpha) = \{K' \in K \bot \alpha : K'' \preceq K' \text{ for all } K'' \in K \bot \alpha\}$$

The relation $\preceq$ “marks off” the most preferred elements of $K \bot \alpha$. When γ is defined in this way, the resulting contraction function is referred to as a *relational partial meet contraction function over K.*

Lemma 2.11 *Let K be a belief set. Any relational partial meet contraction function over K satisfies postulate* (K^-7) *for contraction over K.*

A straightforward extension is to require $\preceq$ be transitive. In this case γ is known as *transitively relational* and the resulting contraction as a *transitively relational partial meet contraction function over K.*

Lemma 2.12 *Let K be a belief set. Any transitively relational partial meet contraction function over K satisfies postulate* (K^-8) *for contraction over K.*

The following theorem supports the utility of such a construction.

Theorem 2.13 *Let K be a belief set and $-$ be a contraction function defined over K. Then $-$ is a transitively relational partial meet contraction function over K* if and only if *it satisfies postulates* (K^-1) — (K^-8) *over K.*

The respective revision operation defined via the Levi identity satisfies postulates (K^*1) — (K^*8). It can also be shown that requiring $\preceq$ to be connected does not lead to further contraction postulates.

---

[18] Cf. the fourth rationality criterion. Each element of $\gamma(K \bot \alpha)$ contains beliefs held in higher regard. The beliefs held in highest regard are those common to the best elements returned of $\gamma(K \bot \alpha)$.

2.2.2 Grove's Sphere Semantics

Grove [19] developed a "sphere semantics" for the AGM framework inspired by Lewis' [32] semantics for counterfactual reasoning.[19] He concentrated on revision functions although the idea is easily extended to deal with contraction (via the Harper Identity) and expansion.

Grove views maximally consistent sets of formulae (consistent complete theories) as "possible worlds". He places an ordering over the set $\mathcal{M}_{\mathcal{L}}$ of all possible worlds. The possible worlds consistent with any set K are denoted $[K]$ and may be defined as follows.

$$[K] = \begin{cases} \{m \in \mathcal{M}_{\mathcal{L}} : K \subseteq m\} & \text{if } K \neq K_{\perp} \\ \emptyset & \text{otherwise} \end{cases}$$

In a similar fashion, the possible worlds consistent with a formula α are denoted $[\alpha]$ and defined as $[\alpha] = [\{\alpha\}]$ (i.e., $[\alpha] = \{m \in \mathcal{M}_{\mathcal{L}} : \alpha \in m\}$). We also define a function $th : 2^{\mathcal{M}_{\mathcal{L}}} \to \mathcal{K}$ mapping sets of possible worlds to belief sets. For any $X \subseteq \mathcal{M}_{\mathcal{L}}$ we have

$$(\text{Def } th) \qquad th(X) = \begin{cases} \bigcap\{m \in X\} & \text{for } X \subseteq \mathcal{M}_{\mathcal{L}} \text{ and } X \neq \emptyset \\ K_{\perp} & \text{if } X = \emptyset \end{cases}$$

We reproduce the following properties, listed by Grove [19], for reference.

Lemma 2.14 *Properties of* th [19].

(i) $th([K]) = K$ *for all belief sets (i.e., theories)* K *if the underlying logic is compact*

(ii) $th(X) \neq K_{\perp}$ *if and only if* X *is nonempty*

(iii) For any sentence $\alpha \in \mathcal{L}$ *and* $X \subseteq \mathcal{M}_{\mathcal{L}}$, $th(X \cap [\alpha]) = Cn(th(X) \cup \{\alpha\})$

(iv) For $X,\ X' \subseteq \mathcal{M}_{\mathcal{L}}$, *if* $X \subseteq X'$, *then* $th(X') \subseteq th(X)$

(v) For $K,\ K' \in \mathcal{K}$, *if* $K \subseteq K'$, *then* $[K'] \subseteq [K]$

A *sphere* is defined to be a set of possible worlds. A *system of spheres centred on* K is an ordering over sets of possible worlds where $[K]$ is the innermost sphere and $\mathcal{M}_{\mathcal{L}}$ the outermost sphere. It can be formally defined as follows.

---

[19] Grove's idea can be viewed as a semantics insofar as it gives a "picture" for AGM belief change. Strictly speaking however, it deals with syntactic objects.

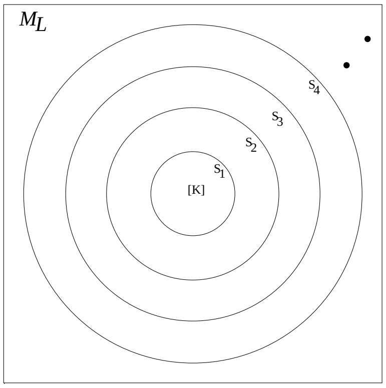

Figure 1: A system of spheres centred on $[K]$.

Definition 2.2 [19]
Let $\mathcal{S}$ be any collection of subsets of $\mathcal{M}_{\mathcal{L}}$. We call $\mathcal{S}$ a system of spheres, centred on X for some subset $X \subseteq \mathcal{M}_{\mathcal{L}}$, if it satisfies the following conditions:

(S1) *$\mathcal{S}$ is totally ordered by $\subseteq$; that is, if $U,\ V \in \mathcal{S}$, then $U \subseteq V$ or $V \subseteq U$*

(S2) *X is the $\subseteq$-minimum of $\mathcal{S}$ (i.e., $X \in \mathcal{S}$ and if $U \in \mathcal{S}$, then $X \subseteq U$)*

(S3) *$\mathcal{M}_{\mathcal{L}}$ is in $\mathcal{S}$ (the largest element of $\mathcal{S}$)*

(S4) *If $\alpha \in \mathcal{L}$, and there is any sphere in $\mathcal{S}$ intersecting $[\alpha]$, then there is a smallest sphere in $\mathcal{S}$ intersecting $[\alpha]$ (there is a sphere $U \in \mathcal{S}$ such that $U \cap [\alpha] \neq \emptyset$, and $V \cap [\alpha] \neq \emptyset$ implies $U \subseteq V$ for all $V \in \mathcal{S}$)*

A pictorial representation of a system of spheres centred on $[K]$ is given in Figure 1.

Condition (S4) guarantees that if any formula α has worlds intersecting $\mathcal{M}_{\mathcal{L}}$ then there is a smallest sphere (in the sense of set inclusion) or innermost sphere in $\mathcal{S}$ intersecting $[\alpha]$. We shall denote such a sphere by $c_{\mathcal{S}}(\alpha)$. If $[\alpha]$ does not intersect any sphere in $\mathcal{S}$ (i.e., $[\alpha] \cap \mathcal{M}_{\mathcal{L}} = \emptyset$), then $c_{\mathcal{S}}(\alpha) = \mathcal{M}_{\mathcal{L}}$ (note that this will only occur whenever $[\alpha] = \emptyset$ by condition (S3)).

With any system of spheres $\mathcal{S}$ centred on $[K]$, we can associate a function $f_{\mathcal{S}} : \mathcal{L} \to 2^{\mathcal{M}_{\mathcal{L}}}$ defined in the following manner for any $\alpha \in \mathcal{L}$

$$(\text{Def } f_{\mathcal{S}}) \qquad\qquad f_{\mathcal{S}}(\alpha) = [\alpha] \cap c_{\mathcal{S}}(\alpha)$$

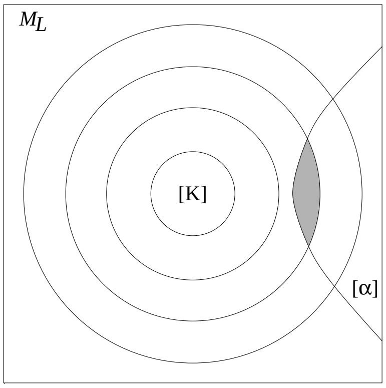

Figure 2: Sphere semantics for belief revision showing $[K^*_\alpha]$ shaded.

Intuitively, the function $f_{\mathcal{S}}$ can be viewed as selecting those α-worlds[20] in $\mathcal{M}_{\mathcal{L}}$ that are "closest" to $[K]$. In other words, it selects the innermost α-worlds.

The sphere semantics for a revision operation is now simply specified as follows.

$$[K^*_\alpha] = f_{\mathcal{S}}(\alpha)$$

That is, the worlds corresponding to a revision of K by α are exactly those α-worlds closest to $[K]$. Such a choice is motivated by the Principle of Minimal Change interpreted with respect to the sphere model outlined above and taking minimality to be "proximity" to $[K]$. It is illustrated in Figure 2 (with $[K^*_\alpha] = f_{\mathcal{S}}(\alpha) = c_{\mathcal{S}}(\alpha) \cap [\alpha]$ shaded).

The following two representation theorems show that the given semantics is appropriate.

Theorem 2.15 *[19]*
Let $\mathcal{S}$ be any system of spheres in $\mathcal{M}_{\mathcal{L}}$ centred on $[K]$ for some belief set K in $\mathcal{K}$. If one defines, for any $\alpha \in \mathcal{L}$, K^*_α to be $th(f_{\mathcal{S}}(\alpha))$, then the postulates (K*1) — (K*8) *are satisfied.*

Theorem 2.16 *[19]*
Let $* : \mathcal{K} \times \mathcal{L} \to \mathcal{K}$ be any function satisfying postulates (K*1) — (K*8)*. Then for any (fixed) belief set K there is a system of spheres on $\mathcal{M}_{\mathcal{L}}$, $\mathcal{S}$ say, centred on $[K]$ and satisfying $K^*_\alpha = th(f_{\mathcal{S}}(\alpha))$, for all $\alpha \in \mathcal{L}$.*

---

[20] An α-world is any world $m \in \mathcal{M}_{\mathcal{L}}$ in which α holds (i.e., $\alpha \in m$). $[\alpha]$ is, of course, the set of all α-worlds.

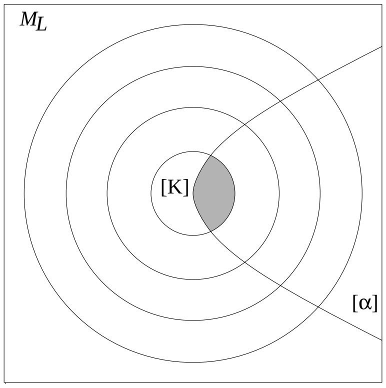

Figure 3: Sphere semantics for belief expansion showing $[K^+_\alpha]$ shaded.

The semantics for belief expansion of an epistemic state K by epistemic input α is now straightforward to determine. In the principal case where $\neg\alpha \notin K$ we have that α is consistent with K and therefore $[K] \cap [\alpha] \neq \emptyset$. That is, the closest α-worlds reside within the innermost sphere $[K]$ and the worlds consistent with the expanded epistemic state are thus given by

$$[K^+_\alpha] = [K] \cap [\alpha]$$

This situation is illustrated in Figure 3. In the case that $\neg\alpha \in K$, we have $K^+_\alpha = K_\perp$. However, in this case $[K] \cap [\alpha] = \emptyset$ and so again, $[K^+_\alpha] = [K] \cap [\alpha]$.

The sphere semantics for belief contraction is slightly more involved though not all that complicated. In fact, it can be easily obtained from that of revision using the Harper Identity. In this situation we are losing information and hence increasing the number of possible worlds. In contracting an epistemic state K by epistemic input α we need to supplement the worlds in $[K]$. Specifically, we must at least incorporate some $\neg\alpha$-worlds otherwise α would still be accepted in the contracted epistemic state and therefore violate the postulate of success for belief contraction.[21] In accordance with the Principle of Minimal Change we should add the closest $\neg\alpha$-worlds. Therefore, the worlds consistent with the new epistemic state may be obtained by

$$[K^-_\alpha] = [K] \cup f_{\mathcal{S}}(\neg\alpha)$$

This situation is illustrated in Figure 4.

---

[21] We are discussing the principal case here in which $\alpha \in K$. If $\alpha \notin K$ then no change in worlds occurs.

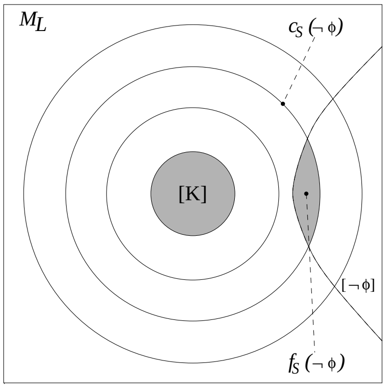

Figure 4: Sphere semantics for belief contraction showing $[K^-_\alpha]$ shaded.

2.2.3 Epistemic Entrenchment

It was shown by Grove [19] that an ordering over possible worlds is equivalent to an ordering over the formulae of $\mathcal{L}$. A more popular treatment along these lines was developed by Gärdenfors and Makinson [14] and is known as *epistemic entrenchment*. Intuitively, such an ordering represents a preference ordering over formulae. Epistemic entrenchment is motivated, to a large extent, by the fourth rationality criterion above.

In contraction, less entrenched formulae would be removed in preference to more deeply entrenched formulae. An epistemic entrenchment ordering may be formally defined as follows.

Definition 2.3 *([14]) An ordering $\leq$ over $\mathcal{L}$ is an epistemic entrenchment ordering if it satisfies the following conditions:*

(SEE1) For any $\alpha,\ \beta,\ \gamma \in \mathcal{L}$, if $\alpha \leq \beta$ and $\beta \leq \gamma$ then $\alpha \leq \gamma$ (transitivity)
(SEE2) For any $\alpha,\ \beta \in \mathcal{L}$, if $\{\alpha\} \vdash \beta$ then $\alpha \leq \beta$ (dominance)
(SEE3) For any $\alpha,\ \beta \in \mathcal{L}$, either $\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$ (conjunctiveness)
(SEE4) When $K \neq K_\perp$, $\alpha \notin K$ iff $\alpha \leq \beta$ for all $\beta \in \mathcal{L}$ (minimality)
(SEE5) If $\beta \leq \alpha$ for all $\beta \in \mathcal{L}$, then $\vdash \alpha$ (maximality)

The first postulate simply states that an epistemic entrenchment ordering is transitive. The Dominance postulate is based on the rationale that, whenever a formula α entails a formula β and one or the other must be given up, a smaller change would result from abandoning α. Giving up β alone is not possible since, being a consequence of α, it would be retained in the resulting belief set. Giving up α alone, on the other hand, may be possible. Therefore, in general, giving up α

would imply a smaller change than giving up β. Hence, β cannot be strictly less entrenched than α. This postulate is clearly motivated by the Principle of Minimal Change. The Conjunctiveness postulate says that removing $\alpha \wedge \beta$ can be accomplished by removing either α or β. The minimality postulate states that non-beliefs are minimally entrenched. The maximality postulate, on the other hand, says that logical truths are maximally entrenched; logical truths are the hardest to give up.[22] Essentially then, an epistemic entrenchment represents a total preorder over the formulae of the language in which tautologies are maximally entrenched and non-beliefs minimally entrenched.

The first three postulates (SEE1) — (SEE3) turn out to be quite significant and any ordering satisfying them is referred to as an *expectations ordering* [15]. Such orderings provide a strong link between the AGM account of belief revision and nonmonotonic inference [15, 35]. Gärdenfors and Makinson supply the following properties satisfied by expectations orderings, some of which will be useful in proving results later in this dissertation.

Lemma 2.17 *(*[14]*)*

(i) $\alpha \leq \beta$ *or* $\beta \leq \alpha$ *(Connectivity)*
(ii) *If* $\beta \wedge \gamma \leq \alpha$*, then* $\beta \leq \alpha$ *or* $\gamma \leq \alpha$
(iii) $\alpha < \beta$ *iff* $\alpha \wedge \beta < \beta$
(iv) *If* $\gamma \leq \alpha$ *and* $\gamma \leq \beta$*, then* $\gamma \leq \alpha \wedge \beta$
(v) *If* $\alpha \leq \beta$*, then* $\alpha \leq \alpha \wedge \beta$

Foo [10] also investigated epistemic entrenchment and provides the following further properties related to expectations orderings and epistemic entrenchment orderings. Note that we may write $\alpha = \beta$ for $\alpha \leq \beta$ *and* $\beta \leq \alpha$. Also, $\alpha < \beta$ is a shorthand for $\alpha \leq \beta$ and $\beta \not\leq \alpha$. We denote the greater of a set of formulae Γ by $max\{\Gamma\}$ and the lesser by $min\{\Gamma\}$. Those properties in the next lemma relate to expectations orderings and some will be helpful later on.

Lemma 2.18 *(*[10]*)*

(i) *If* $\alpha < \beta$ *and* $\beta < \gamma$*, then* $\alpha < \gamma$
(ii) *If* $\alpha \leq \beta$ *and* $\beta < \gamma$*, then* $\alpha < \gamma$
(iii) *If* $\beta < \gamma$ *and* $\beta < \alpha$*, then* $\beta < \alpha \wedge \gamma$
(iv) $\alpha \not< \alpha$ *for any* $\alpha \in \mathcal{L}$
(v) *If* $\alpha < \beta$*, then* $\alpha < \gamma \vee \beta$ *for any* $\gamma \in \mathcal{L}$
(vi) *If* $\alpha < \beta$*, then* $\alpha < \alpha \vee \beta$

---

[22] In fact, they cannot be given up at all given our second rationality criterion.

(vii) *If* $\alpha < \beta$*, then* $\alpha \wedge \gamma < \beta$ *for any* $\gamma \in \mathcal{L}$
(viii) *If* $\alpha \leq \beta$*, then* $\alpha \wedge \gamma \leq \beta$
(ix) *If* $\beta \wedge \gamma < \alpha$*, then* $\beta < \alpha$ *or* $\gamma < \alpha$
(x) *If* $\alpha \leq \beta$*, then* $\alpha \wedge \gamma \leq \beta \wedge \gamma$
(xi) *If* $\alpha = \beta$*, then* $\alpha \wedge \gamma = \beta \wedge \gamma$
(xii) $\alpha \wedge \beta = min\{\alpha,\ \beta\}$
(xiii) $\alpha \vee \beta \geq max\{\alpha,\ \beta\}$
(xiv) *If* $\alpha = \alpha \vee \beta$*, then* $max\{\alpha,\ \beta\} = \alpha$
(xv) $\alpha = \alpha \vee \beta$ *or* $\beta = \alpha \vee \beta$ *iff* $max\{\alpha,\ \beta\} = \alpha \vee \beta$
(xvi) α *and* β *are not independent*[23] *iff* $\alpha = \alpha \vee \beta$ *or* $\beta = \alpha \vee \beta$

The following results relate to epistemic entrenchment orderings.

Lemma 2.19 *(*[10]*)*

(i) *If* $\nvdash \alpha$ *and* $\vdash \beta$*, then* $\alpha < \beta$
(ii)[24] *Suppose* $\beta \in K$*. If* $\alpha < \beta$*, then* $\beta \in K^{-}_{\alpha}$
(iii) *For all* α *and* $\beta \notin K$*,* $\alpha \leq \beta$ *and* $\beta \leq \alpha$
(iv) $\alpha \notin K$ *iff* $\alpha < \beta$ *for all* $\beta \in K$

An epistemic entrenchment ordering $\leq$ for a particular belief set K may be constructed from a contraction function $-$ using the following condition.

$$(\mathrm{C}\ \leq) \qquad \alpha \leq_{-} \beta \text{ iff } \alpha \notin K^{-}_{\alpha\wedge\beta} \text{ or } \vdash \alpha \wedge \beta$$

The principal part of the condition states that β is at least as epistemically entrenched as α whenever α is removed from K in contracting K by $\alpha \wedge \beta$ since, to contract by $\alpha \wedge \beta$, only one of α or β need be removed and the fact that α has been removed means that it cannot be strictly more entrenched than β (otherwise, only β need be given up). In the case where α and β are both tautological, they are equally entrenched.

More importantly, it is possible to construct a contraction function $-_{\leq}$ (restricted to a particular K)[25] from an epistemic entrenchment ordering as follows.

$$(\mathrm{C}\ -) \qquad \beta \in K^{-\leq}_{\alpha} \text{ iff both } \beta \in K \text{ and either } \alpha < \alpha \vee \beta \text{ or } \vdash \alpha$$

---

[23] Two formulae are independent if one can be removed without affecting the other.

[24] The proof of this property requires condition (C$-$) which we shall introduce shortly.

[25] Note that given an epistemic entrenchment relation $\leq$, the belief set over which it is restricted is easily determined as $K = \{\alpha : \beta < \alpha \text{ for some } \beta \in \mathcal{L}\}$.

Clearly, any formula not in the original epistemic state is not going to occur in the contracted epistemic state. In the situation where the epistemic input is a logical truth, it cannot be retracted and therefore no change is made. Otherwise, we note that by the recovery postulate, $\neg\alpha \vee \beta \in K^-_\alpha$ for any belief $\beta \in K$. Now, if the disjunction $\alpha \vee \beta$ of the epistemic input and some belief β is more entrenched than the epistemic input itself, then this disjunction is going to be retained. These two facts imply that β will remain in the contracted state. The following representation theorems show the appropriateness of the epistemic entrenchment ordering and the conditions given above.

Theorem 2.20 *[14]*
Let $K \in \mathcal{K}$ be a belief set and $\leq$ be an epistemic entrenchment over K. If for any $\alpha \in \mathcal{L}$, we define K^-_α using $(\mathrm{C}\ -)$*, then* (K^-1) — (K^-8) *are satisfied as well as the condition* $(\mathrm{C} \leq)$.

Theorem 2.21 *[14]*
Let $- : \mathcal{K} \times \mathcal{L} \rightarrow \mathcal{K}$ be any function satisfying (K^-1) — (K^-8)*. Then, for any belief set $K \in \mathcal{K}$, if we define $\leq$ using condition* $(\mathrm{C} \leq)$*, then $\leq$ is an epistemic entrenchment ordering (i.e., it satisfies* (SEE1) — (SEE5)*) and also satisfies condition* $(\mathrm{C}\ -)$.

These three constructions are arguably the most important for the AGM framework. They are clearly related as evidenced by the representation theorems. The interested reader is referred to Gärdenfors [12] and Peppas and Williams [45] for a further discussion of these relationships.

2.2.4 Safe Contraction

The construction termed *safe contraction* [2, 3] combines, in a certain sense, elements common to both epistemic entrenchment orderings and partial meet contraction functions. On the one hand, it is assumed that an acyclic relation $<$ of the elements of K is given.[26] Moreover, we consider the minimal subsets of an epistemic state K that imply epistemic input α (in partial meet contraction functions, however, note that we deal with maximal subsets of K failing to imply α). Such a subset may be defined as follows.

Definition 2.4 *A set K' is a minimal subset of K implying α if and only if*

---

[26] Alchourrón and Makinson refer to an acyclic relation as a “hierarchy”. It will be irreflexive and asymmetric.

(i) $K' \subseteq K$

(ii) $K' \vdash \alpha$

(iii) $K'' \nvdash \alpha$ *for any* $K'' \subset K'$

The set of all minimal subsets of K implying α is denoted $K \Downarrow \alpha$ (cf. $K \bot \alpha$).

Definition 2.5 *Any belief* $\beta \in K$ *is said to be* safe *with respect to* α *if and only if* β *is not minimal under* $<$ *with respect to the elements of any* $K' \in K \Downarrow \alpha$*. The set of all safe elements of K is denoted* $K \setminus \alpha$.

A belief is safe if it is not "culpable" for the presence of α. Intuitively, one element must be removed from each subset of K in $K \Downarrow \alpha$. The ordering $<$ helps us choose which element to remove from each subset. The remaining beliefs are safe and can be used to determine the *safe contraction* of a belief set K by α (modulo $<$). Specifically, we define $K^-_\alpha = Cn(K \setminus \alpha)$.

Safe contraction functions satisfy the six basic postulates for contraction. It is interesting to investigate particular types of hierarchies $<$ over beliefs.

Definition 2.6 *If* K *is a belief set and* $<$ *is a hierarchy then, for all* $\alpha,\ \beta,\ \gamma \in K$

(i) $<$ continues up $\vdash$ over K *if and only if* $\alpha < \beta$ *and* $\beta \vdash \gamma$ *imply* $\alpha < \gamma$

(ii) $<$ continues down $\vdash$ over K *if and only if* $\alpha \vdash \beta$ *and* $\beta < \gamma$ *imply* $\alpha < \gamma$

(iii) $<$ is virtually connected over K *if and only if* $\alpha < \beta$ *implies either* $\alpha < \gamma$ *or* $\gamma < \beta$

It can be shown [2] that, if $<$ continues up or down $\vdash$, then the resulting safe contraction function satisfies the postulate of intersection (K$^-$7) and, if $<$ is virtually connected over K, it satisfies the postulate of conjunction (K$^-$8) over K. The following representation theorem holds at least when K consists of a finite number of logically equivalent sentences (i.e., when K is partitioned into a finite number of equivalence classes by the consequence relation $\vdash$).

Theorem 2.22 *Let* K *be a belief set. A safe contraction function* $-$ *is generated by a hierarchy* $<$ *that continues up and down* $\vdash$ *over* K *and is virtually connected if and only if* $-$ *is a transitively relational partial meet contraction function over* K.

3 Spohn — Ordinal Conditional Functions

An alternative approach to the problem of belief revision has been proposed by Spohn [52]. We shall only give a brief outline here as we do not make much use of this approach in this dissertation.

Spohn bases his account on possible worlds although one need not identify these with the possible worlds considered by Grove [19][27]. The set of all possible worlds is denoted W. Possible worlds are considered to be ordered by a *grading of disbelief*. An *ordinal conditional function* k is used for this purpose, assigning an ordinal to each world $w \in W$. The smaller the ordinal assigned to a world, the more plausible (less disbelieved) it is (0 being the smallest ordinal). In this manner, k specifies an epistemic state.

Since propositions can be identified with sets of worlds, it is also possible to talk about the grading of disbelief of a proposition α. It is simply that of the most plausible of its α-worlds i.e., $k(\alpha) = min\{k(w) : \alpha \in w\}$. A grading of disbelief possesses two important properties

- Either $k(\alpha) = 0$ or $k(\neg\alpha) = 0$ for any proposition α
- $k(\alpha \cup \beta) = min\{k(\alpha),\ k(\beta)\}$ for consistent α and β.

We can therefore say that a proposition α is believed (or accepted) in an epistemic state induced by k if and only if $k(\neg\alpha) > 0$. This proposition α is said to be believed with *degree of firmness* $k(\neg\alpha)$. An ordinal conditional function then, allows us to say whether one proposition is more firmly believed (more plausible) than another proposition β in an epistemic state. In this way, the number of possible epistemic attitudes is greater than those possible with the AGM and therefore ordinal conditional functions are more discriminating.

Another important difference between Spohn's framework and the AGM is the manner in which belief change is effected. Spohn takes epistemic inputs to consist not only of a proposition α but also of an ordinal. Intuitively, the ordinal represents the degree of firmness α should acquire after the change takes place. This means that belief expansion, contraction and revision, as we have come to know them in the AGM, can all be captured by a single mechanism. The actual belief change process is known as *conditionalisation*. The basic idea is that, for epistemic input α (proposition) and i (ordinal) the α and $\neg\alpha$ worlds are "shifted" relative to each other in order to assign α degree of firmness i. An example is illustrated in Figure 5. In the ordinal conditional function on the left $k(\alpha) = 1$ and $k(\neg\alpha) = 0$ (i.e., $\neg\alpha$ is believed with degree of firmness 1) while that on the right

---

[27] Spohn's possible worlds can be thought of as uninterpreted points.

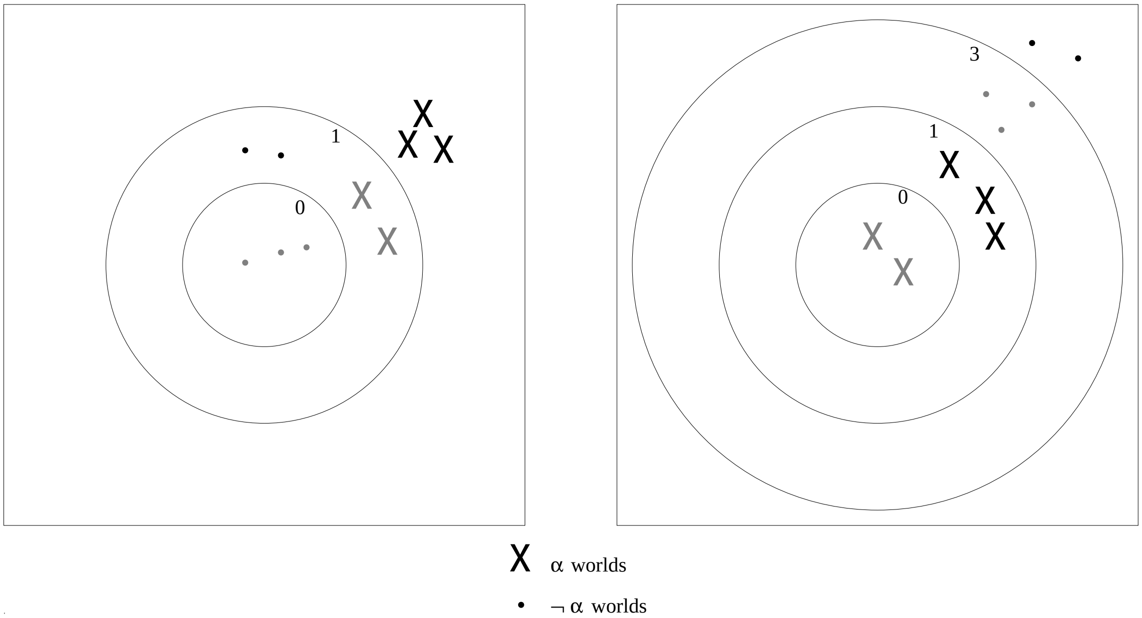

Figure 5: An example of Spohn's approach to belief revision.

shows the result of conditionalisation on input $\alpha, 3$. Another advantage of Spohn's approach, evidenced by this example is that it permits *iterated revision*. That is, it is possible to perform a sequence of belief changes due to the fact that the ordering on worlds (grading of disbelief) is still defined after every change. It is not, at first, clear how this is achievable in the AGM since there is no selective mechanism (i.e., system of spheres, epistemic entrenchment, etc.) defined after a change. However, a number of authors have attacked the problem [18, 29, 38, 39, 50, 53].

4 Summary and Discussion

Belief revision is the study of the dynamics of epistemic states. The two main approaches to modelling epistemic states are known as foundationalism and coherentism. Essentially, foundationalism posits the existence of a select set of epistemologically basic beliefs whereas coherentism denies the existence of any such beliefs at all. Formally however, any such difference has been called into question. Dixon and Foo [6] show, in the case of contraction, how ATMS behaviour can be achieved through a particular ordering of beliefs in an epistemic entrenchment ordering. Only the relative ordering of certain formulae need be specified, giving rise to a *partial epistemic entrenchment ordering*. This ordering characterises a class of epistemic entrenchment orderings, any of which exhibit the same contraction behaviour as a particular ATMS context. Del Val [5] goes even further, showing that, for a finite propositional language, a mathematical definition of a coherence revision operator based on Katsuno and Mendelzon's [26] version of the AGM and a definition of a foundational revision operator motivated by syntax-based approaches to belief revision (see [33, 40], for example) lead to identical classes of revision operators. This result, however, only shows the equivalence of operators satisfying the definitions given and leaves open the connection between coherence and foundational theories in general. Moreover, one must keep in mind that these theories concern the nature of epistemic states not the method employed to move from one epistemic state to another. We shall stick with the more intuitive descriptions of the theories given here. Arguments for and against both theories can be found in Gärdenfors [13] and Doyle [8].

Our main concern here is with the (purportedly) coherent AGM framework, due principally to its well developed logical theory. Katsuno and Mendelzon [27] claim that the AGM is well suited to situations in which an agent is reasoning about a static world but does not have full information about it. They offer an account of an alternative belief change operator, known as *belief update*, claimed to be suited to reasoning about dynamic worlds. Peppas [44] investigates the relationship between the two approaches but since it is not central to our concerns here, we shall not consider it further.

References

[1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985. Classic AGM paper.

[2] Carlos E. Alchourrón and David Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44:405–422, 1985.

[3] Carlos E. Alchourrón and David Makinson. Maps between some different kinds of contraction function: The finite case. *Studia Logica*, 45:187–198, 1986.

[4] Johan de Kleer. An assumption-based TMS. *Artificial intelligence*, 32:127–162, 1986.

[5] Alvaro del Val. On the relation between the coherence and foundations theories of belief revision. In *Proceedings of the Thirteenth Joint Conference on Artificial Intelligence*, pages 909–914, 1993.

[6] Simon Edmund Dixon and Norman Y. Foo. Connections between the ATMS and AGM belief revision. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 534–539, 1993.

[7] Jon Doyle. A truth maintenance system. *Artificial Intelligence*, 12:231–272, 1979.

[8] Jon Doyle. Reason maintenance and belief revision: Foundations versus coherence theories. In Peter Gärdenfors, editor, *Belief Revision*, pages 29–51. Cambridge University Press, 1992.

[9] Charles Elkan. A rational reconstruction of non-monotonic truth maintenance systems. *Artificial Intelligence*, 43:219–234, 1990.

[10] Norman Y. Foo. Observations on AGM entrenchment. Technical Report 389, Department of Computer Science, University of Sydney, June, 1990.

[11] André Fuhrmann. Theory contraction through base revision. *Journal of Philosophical Logic*, 20:175–203, 1991.

[12] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge Massachusetts, 1988. Classic introduction to the AGM theory of belief change. Becoming a bit dated but still worth reading.

[13] Peter Gärdenfors. The dynamics of belief systems: Foundations vs. coherence theories. *Revue Internationale de Philosophie*, 44:24–46, 1990. Puts forward the premise that AGM belief change corresponds to a coherence theory.

[14] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspect of Reasoning About Knowledge*, pages 83–96, 1988. Contains the main results regarding the epistemic entrenchment construction for AGM contraction functions.

[15] Peter Gärdenfors and David Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65:197–245, 1994.

[16] Peter Gärdenfors and Hans Rott. Belief revision. In Dov M. Gabbay, Christopher J. Hogger, and John A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming Volume IV: Epistemic and Temporal Reasoning*, pages 35–132. Oxford University Press, 1995. A good introduction to the AGM approach which also stresses belief bases.

[17] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programs. In *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning*, pages 1070–1080, 1988.

[18] Aditya Ghose, Pablo O. Hadjinian, Abdul Sattar, Jia-H. You, and Randy Goebel. Iterated belief change. In *Proceedings of the Sixth Australian Joint Conference Artificial Intelligence*, pages 39–44, 1993.

[19] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988. Details the sphere constructive modelling of AGM revision functions.

[20] Sven Ove Hansson. Belief contraction without recovery. *Studia Logica*, pages 251–260, June 1991.

[21] Sven Ove Hansson. In defense of base contraction. *Synthese*, 91:239–245, 1992.

[22] Sven Ove Hansson. A textbook of belief dynamics: Theory change and database updating. Kluwer, to appear. An excellent comprehensive reference to the AGM approach containing many technical results together with their proofs, 1996.

[23] Sven Ove Hansson and André Fuhrmann. A survey of multiple contractions. *Journal of Logic Language and Information*, 3(1):39–76, 1994.

[24] Sven Ove Hansson. Revision of belief sets and belief bases. In D. Gabbay and Smets, editors, *DRUMS Handbook, Volume 3*, 1998, to appear.

[25] Gilbert H. Harman. *Change in View: Principles of Reasoning*. MIT Press, 1986.

[26] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

[27] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In Peter Gärdenfors, editor, *Belief*

Revision, pages 183–203. Cambridge University Press, 1992. A preliminary version of a paper bearing the same title appearing in James Allen, Richard Fikes and Erik Sandewall (editors), *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, Morgan Kaufmann, CA, pp. 387-394, 1991. Outlines belief update as a different type of belief change.

[28] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence,* 44:167–207, 1990.

[29] Daniel Lehmann. Belief revision, revised. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1534–1540. Montreal, 1995.

[30] Isaac Levi. Subjunctives, dispositions and chances. *Synthese*, 34:423–455, 1977.

[31] Isaac Levi. *The Fixation of Belief and its Undoing*. Cambridge University Press, 1991.

[32] David Lewis. *Counterfactuals*. Harvard University Press, Cambridge, Massachusetts, 1973.

[33] David Makinson. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthese,* 32:347–363, 1985.

[34] David Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16:383–394, 1987.

[35] David Makinson and Peter Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In A. Fuhrmann and M. Morreau, editors, *The Logic Of Theory Change*, pages 185–205. Lecture Notes in Artificial Intelligence 465, Springer Verlag, 1990.

[36] Robert C. Moore. Semantical considerations of nonmonotonic logic. *Artificial Intelligence,* 25:75–94, 1985.

[37] Abhaya C. Nayak. Foundational belief change. *Journal of Philosophical Logic*, 23:495–533, 1994.

[38] Abhaya C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis,* 41:353–390, 1994.

[39] Abhaya C. Nayak, Norman Y. Foo, Maurice Pagnucco, and Abdul Sattar. Changing conditional beliefs unconditionally. In Johan van Benthem and Yoav Shoham, editors, *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge*, pages 119–135, De Zeeuwse Stromen, The Netherlands, 1996.

[40] Bernhard Nebel. Belief revision and default reasoning: Syntax based approaches. In Peter Gärdenfors, editor, *Belief Revision*, pages 52–88. Cambridge University Press, 1992. A revised and extended version of a paper bearing the same title appearing in James Allen, Richard Fikes and Erik Sandewall (editors), *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, Morgan Kaufmann, CA, pp. 417-428, 1991.

[41] Otto Neurath. Protokollsätze. *Erkenntnis*, 3:204–214, 1932. English translation appears in [42].

[42] Otto Neurath. *Philosophical Papers 1913–1946*. D. Reidel Publishing Company, 1983. Edited and Translated by Robert S. Cohen and Marie Neurath with the editorial assistance of Carolyn R. Fawcett.

[43] Erik J. Olsson. *Coherence: Studies in Epistemology and belief Revision*. Acta Universitatis Upsaliensis. *Comprehensive Summaries of Uppsala Dissertations from the Faculty of Arts 2*, 1997.

[44] Pavlos Peppas. *Belief Change and Reasoning about Action — An Axiomatic Approach to Modelling Dynamic Worlds and the Connection to the Logic of Theory Change*. PhD thesis, University of Sydney, 1993.

[45] Pavlos Peppas and Mary-Anne Williams. Constructive modelings for theory change. *Notre Dame Journal of Formal Logic*, 36(1):120–133, Winter 1995.

[46] John L. Pollock. *Contemporary Theories of Knowledge*. Rowan and Littlefield Publishers, 1986.

[47] M. Reinfrank. *Fundamentals and Logical Foundations of Truth Maintenance*. PhD thesis, Department of Computer and Information Science, University of Linkoeping, 1989.

[48] M. Reinfrank, O. Dessler, and G. Brewka. On the relation between truth maintenance and autoepistemic logic. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1206–1212, 1989.

[49] Raymond Reiter and Johan de Kleer. Foundations of assumption-based truth maintenance systems: Preliminary report. In *Proceedings of the National Conference in Artificial Intelligence*, pages 183–188, 1987.

[50] Hans Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic Language and Information*, 1:45–78, 1992.

[51] Ernest Sosa. The raft and the pyramid: Coherence versus foundations in the theory of knowledge. In *Midwest Studies in Philosophy 5*, pages 3–25. University of Minnesota, Morris, 1980.

[52] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper and Brian Skryms, editors, *Causation in Decision, Belief Change, and Statistics, II*, pages 105–134. Kluwer Academic Publishers, 1988.

[53] Mary-Anne Williams, Maurice Pagnucco, Norman Y Foo, and Brailey Sims. Determining explanations using knowledge transmutations. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 822–830. Montreal, 1995.